SimRank Computation on Uncertain Graphs

Rong Zhu, Zhaonian Zou, Jianzhong Li
Department of Computer Science and Technology, Harbin Institute of Technology, China
{rzhu,znzou,lijzh}@hit.edu.cn

arXiv:1512.02714v1 [cs.DB] 9 Dec 2015

***Abstract*—SimRank is a similarity measure between vertices in a graph, which has become a fundamental technique in graph analytics. Recently, many algorithms have been proposed for efficient evaluation of SimRank similarities. However, the existing SimRank computation algorithms either overlook uncertainty in graph structures or is based on an unreasonable assumption (Du et al). In this paper, we study SimRank similarities on uncertain graphs based on the possible world model of uncertain graphs. Following the random-walk-based formulation of SimRank on deterministic graphs and the possible worlds model of uncertain graphs, we define random walks on uncertain graphs for the first time and show that our definition of random walks satisfies Markov's property. We formulate the SimRank measure based on random walks on uncertain graphs. We discover a critical difference between random walks on uncertain graphs and random walks on deterministic graphs, which makes all existing SimRank computation algorithms on deterministic graphs inapplicable to uncertain graphs. To efficiently compute SimRank similarities, we propose three algorithms, namely the baseline algorithm with high accuracy, the sampling algorithm with high efficiency, and the two-phase algorithm with comparable efficiency as the sampling algorithm and about an order of magnitude smaller relative error than the sampling algorithm. The extensive experiments and case studies verify the effectiveness of our SimRank measure and the efficiency of our SimRank computation algorithms.**

I. Introduction

Complicated relationships between entities are often represented by a graph. The similarities between entities can be revealed by analyzing the links between the vertices in a graph. Recently, evaluating similarities between vertices has become a fundamental issue in graph analytics. It plays an important role in many applications, including entity resolution [4], recommender systems [9] and spams detection [3]. Assessing similarities between vertices is also a cornerstone of many graph mining tasks, such as graph clustering [43], frequent subgraph mining [42] and dense subgraph discovery [44].

A lot of similarity measures have been proposed, e.g., Jaccard similarity [13], Dice similarity [6] and cosine similarity [2], which are motivated by the intuition that two vertices are more similar if they share more common neighbors. However, these measures cannot evaluate similarities between vertices with no common neighbors. To address this problem, Jeh and Widom [14] proposed a versatile similarity measure called *SimRank* based on the intuition that *two vertices are similar if their in-neighbors are similar too*. Since SimRank captures the topology of the whole graph, it can be used to assess the similarity between two vertices regardless if they have common neighbors. Hence, SimRank has been widely used over the last decade. A lot of studies [8], [10], [19]–[21], [24], [31], [33], [37], [39] have been done on efficient SimRank computations.

Almost all the studies on SimRank focus on deterministic graphs. However, in recent years, people have realized that uncertainty is intrinsic in graph structures, e.g., protein-protein interaction (PPI) networks. A graph inherently accompanied with uncertainty is called an *uncertain graph*. Considerable researches on managing and mining uncertain graphs [7], [16], [18], [30], [44], [46] have shown that the effects of uncertainty on the quality of results have been undervalued in the past. To the best of our knowledge, the only work of SimRank computation on uncertain graphs has been carried out by Du et al. [7]. Whereas, SimRank on uncertain graphs is important in many applications. We show two examples as follows.

Application 1 (Detecting Similar Proteins). Finding proteins with similar biological functions is of great significance in biology and pharmacy [27], [34]. Traditionally, the similarity between proteins are measured by matching their corresponding DNA's [34]. However, similar DNA sequences may not generate proteins with similar functions. Recent approaches are based on protein-protein interaction (PPI) networks. A PPI network represents interactions between proteins detected by high-throughput experiments. A PPI network reflects functional relationships among proteins more directly. A pair of proteins with high structural-context similarity in a PPI network are more likely to have similar biological functions. However, due to errors and noise in high-throughput biological experiments, uncertainty is inherent in a PPI network. This motivates us to sutdy SimRank similarities on uncertain graphs.

Application 2 (Entity Resolution). Entity Resolution (ER) is a primitive operation in data cleaning. The goal of ER is to find records that refer to the same real-world entity from heterogeneous data sources [4]. Considerable ER algorithms [4], [22], [35] fall into the category of organizing data records as a graph, where vertices represent data records, and edges between records are associated with similarity values. Such graph is typically an uncertain graph since the weights are often normalized into $[0, 1]$ and regarded as probabilities. In the existing graph-based ER algorithms, they aggregate similar vertices into an entity but ignore uncertainty information. For example, the EIF algorithm [22] discards the edges whose weights are less than a threshold and aggregates similar records according to the Jaccard similarity. To take uncertainty into account, we study SimRank similarities on uncertain graphs.

Challenges. The challenges of SimRank on uncertain graphs come from two aspects, namely its formulation and computation. Notice that the SimRank on a deterministic graph can be formulated in the language of *random walks on graphs* [14]. Specifically, the SimRank matrix (i.e., the matrix of SimRank similarities between all pairs of vertices) can be formulated as a (nonlinear) combination of the one-step transition probability matrix (i.e., the matrix of one-step transition probabilities between all pairs of vertices). However, such formulation cannot be adapted to uncertain graphs. This is because, for a deterministic graph, the k-step transition probability matrix $\mathbf{W}^{(k)}$ equals the kth power of the one-step transition probability matrix $\mathbf{W}^{(1)}$, that is, $\mathbf{W}^{(k)} = (\mathbf{W}^{(1)})^k$. However,

as analyzed in this paper, for an uncertain graph, the k-step transition probability matrix $\mathcal{W}^{(k)}$ is *unequal* to the kth power of the one-step transition probability matrix $\mathcal{W}^{(1)}$, i.e., $\mathcal{W}^{(k)} \neq (\mathcal{W}^{(1)})^k$. Unfortunately, the only work of SimRank on uncertain graphs [7] does not solve this problem because it makes an unreasonable assumption that $\mathcal{W}^{(k)} = (\mathcal{W}^{(1)})^k$ for all $k \geq 1$. Hence, the first two challenges are as follows.

1: How to define random walks on uncertain graphs?
2: How to define SimRank on uncertain graphs based on random walks on uncertain graphs?

For a deterministic graph, the SimRank matrix can be approximated using many methods [24], [33], [36], [37], [39], [40]. All these methods are based on the fact that the SimRank matrix is a (nonlinear) combination of the one-step transition probability matrix $\mathbf{W}^{(1)}$. Therefore, the central operations involved in these methods are matrix multiplications with the columns of $\mathbf{W}^{(1)}$ and $(\mathbf{W}^{(1)})^T$. However, all these methods cannot be adapted to compute the SimRank matrix for an uncertain graph because the k-step transition probability matrix $\mathcal{W}^{(k)}$ on uncertain graphs does not satisfy that $\mathcal{W}^{(k)} = (\mathcal{W}^{(1)})^k$. In fact, the SimRank matrix for an uncertain graph is a combination of all transition probability matrices $\mathcal{W}^{(k)}$ for $k = 1, 2, \ldots$. Hence, the challenges in computations are as follows.

3: How to efficiently compute the k-step transition probability matrix $\mathcal{W}^{(k)}$ for an uncertain graph?
4: How to efficiently approximate the SimRank matrix for an uncertain graph?

To deal with the challenges C1–C4 listed above, we study the theory and algorithms on SimRank on uncertain graphs. The studies in this paper are strictly based on the *possible world model* of uncertain graphs [16], [18], [30], [45]. In the possible world model, an uncertain graph represents a probability distribution over all the possible worlds of the uncertain graph. Each possible world is a deterministic graph that the uncertain graph could possibly be in practice. The main contributions of this paper are as follows.

Contribution 1. To the best of knowledge, we are the first to formulate random walks on uncertain graphs totally following the possible world model. We define the k-step transition probability from a vertex u to a vertex v as the probability that a walk stays at u at time n and arrives at v at time $n+k$ in a randomly selected possible world. Our definition satisfies *Markov's property*, that is, for all $n \geq 0$ and all vertices v, the probability that a walk stays at v at time $n+1$ is only determined by the vertex at which the walk stays at time n, independent of all the vertices that the walk has visited at time $0, 1, 2, \ldots, n-1$. One of our main findings is that, for an uncertain graph, the k-step transition probability matrix $\mathcal{W}^{(k)}$ is *not equal to* the kth power of the one-step transition probability matrix $\mathcal{W}^{(1)}$. In case when there is no uncertainty involved in graphs, our definition of random walks on uncertain graphs degenerates to random walks on deterministic graphs.

Contribution 2. Based on the model of random walks on uncertain graphs, we define the SimRank measure on uncertain graphs. The SimRank similarity between two vertices u and v is formulated as the combination of the probabilities that two random walks starting from u and v, respectively, meet at the same vertex after k transitions for all $k = 1, 2, \ldots$. Since $\mathcal{W}^{(k)} \neq (\mathcal{W}^{(1)})^k$, we cannot formulate the SimRank matrix in a recursive form of $\mathcal{W}^{(1)}$ only. Thus, the existing algorithms for SimRank computations cannot be used to evaluate SimRank similarities on uncertain graphs.

Contribution 3. We propose three algorithms for approximating the SimRank similarity between two vertices. The central idea of these algorithms is approximating the SimRank similarity between two vertices u and v by combining the probabilities that two random walks starting from u and v, respectively, meet at the same vertex after k transitions for $1 \leq k \leq n$, where n is a sufficiently large number. We prove that the approximate value converges to the exact value as $n \to +\infty$. Moreover, the approximation error exponentially decreases as n becomes larger.

The three SimRank computation algorithms proposed in this paper adopt different approaches to computing transition probability matrices. The first algorithm exactly computes the transition probability matrices $\mathcal{W}^{(1)}, \mathcal{W}^{(2)}, \ldots, \mathcal{W}^{(n)}$. The second algorithm approximates $\mathcal{W}^{(1)}, \mathcal{W}^{(2)}, \ldots, \mathcal{W}^{(n)}$ via sampling. To make a tradeoff between efficiency and accuracy, we propose the third algorithm called the two-phase algorithm, which works in two phases. Let $1 \leq l \leq n$. In the first phase, we exactly compute $\mathcal{W}^{(1)}, \mathcal{W}^{(2)}, \ldots, \mathcal{W}^{(l)}$; in the second phase, we approximate $\mathcal{W}^{(l+1)}, \mathcal{W}^{(l+2)}, \ldots, \mathcal{W}^{(n)}$ by sampling. Finally, we combine these results to approximate the SimRank similarities. By carefully selecting l, the two-phase algorithm can achieve comparable efficiency as the sampling algorithm and about an order of magnitude smaller relative error than the sampling algorithm. Furthermore, we develop a new technique to share the common steps within a large number of independent sampling processes, which decreases the total sampling time by 1–2 orders of magnitude.

Contribution 4. We conducted extensive experiments on a variety of uncertain graph datasets to evaluate our proposed algorithms. The experimental results verify both the effectiveness and the convergence of our SimRank measure. The two-stage algorithm is much more efficient than the baseline algorithm on large uncertain graphs, and its relative error is about an order of magnitude smaller than the sampling algorithm. Moreover, our speeding-up technique can make the sampling process 1–2 orders of magnitude faster without harming the relative errors of the results. We also performed two interesting case studies on detecting similar proteins and entity resolution as we stated above. The results verify the effectiveness of our SimRank similarity measure.

The rest of this paper is organized as follows. Section 2 reviews some preliminaries. Section 3 gives a formal definition of random walks on uncertain graphs. Section 4 proposes the algorithm for computing the k-step transition probability matrices of an uncertain graph. Section 5 formulates the SimRank measure on uncertain graphs. Section 6 proposes three SimRank computation algorithms and the speeding-up technique. Section 7 reports the extensive experimental results. Section 8 overviews the related work. Finally, this paper is concluded in Section 9.

II. Preliminaries

In this section we review some preliminary knowledge, including *random walks on graphs*, the *SimRank similarity measure* and the model of *uncertain graphs*.

Random Walks on Graphs. A (deterministic) *graph* is a pair (V, E), where V is a set of *vertices*, and $E \subseteq V \times V$. Each

element $(u, v) \in E$ is said to be an *arc* connecting vertex u to vertex v. In this paper we consider *directed* graphs, in which (u, v) and (v, u) refer to different arcs.

Let G be a directed graph. We use $V(G)$ and $E(G)$ to denote the vertex set and the arc set of G, respectively. A vertex u is said to be an *in-neighbor* of a vertex v if (u, v) is an arc. Meanwhile, v is an *out-neighbor* of u. Let $I_G(v)$ and $O_G(v)$ denote the sets of in-neighbors and out-neighbors of a vertex v in G, respectively. A *walk* on G is a sequence of vertices $W = v_0, v_1, \ldots, v_n$ such that (v_i, v_{i+1}) is an arc for $0 \leq i \leq n-1$. The *length* of W, denoted by $|W|$, is n. A sequence of random variables $X_0, X_1, X_2, \ldots$ over $V(G)$ is a *random walk* on G if it satisfies *Markov's property*, that is,

$$\begin{aligned}&\Pr(X_i = v_i | X_0 = v_0, X_1 = v_1, \ldots, X_{i-1} = v_{i-1})\\ =&\Pr(X_i = v_i | X_{i-1} = v_{i-1})\end{aligned}$$

for all $i \geq 1$ and all $v_0, v_1, \ldots, v_i \in V(G)$. For any $u, v \in V(G)$, $\Pr(X_i = v | X_{i-1} = u)$ represents the probability that the random walk, when on vertex u at time $i-1$, will next make a transition onto vertex v at time i. Particularly,

$$\Pr(X_i = v | X_{i-1} = u) = \begin{cases} \frac{1}{|O_G(u)|} & \text{if } (u, v) \in E(G),\\ 0 & \text{otherwise.}\end{cases}$$

Note that $\Pr(X_i = v | X_{i-1} = u)$ is fixed for all $i \geq 1$, so we denote the value of $\Pr(X_i = v | X_{i-1} = u)$ by $\Pr(u \to_1 v)$, which is called the *one-step transition probability* from vertex u to vertex v. Therefore, for all $v_0, v_1, \ldots, v_n \in V(G)$, the probability that a random walk $X_0, X_1, X_2, \ldots$, when starting from vertex v_0 at time 0, will later be on vertex v_i at time i for $1 \leq i \leq n$ is

$$\begin{aligned}&\Pr(X_1 = v_1, X_2 = v_2, \ldots, X_n = v_n | X_0 = v_0)\\ =&\prod_{i=1}^{n} \Pr(v_{i-1} \to_1 v_i).\end{aligned} \tag{1}$$

For all $i \geq 0$, $k \geq 1$, and all $u, v \in V(G)$, the probability that a random walk on vertex u at time i will later be on vertex v after k additional transitions is

$$\begin{aligned}&\Pr(X_{i+k} = v | X_i = u)\\ =&\sum_{v_1, v_2, \ldots, v_{k-1} \in V(G)} \Pr(X_{i+1} = v_1, \ldots, X_{i+k-1} = v_{k-1}, X_{i+k} = v | X_i = u),\end{aligned}$$

By Eq. (1), $\Pr(X_{i+k} = v | X_i = u)$ is fixed for all $i \geq 0$, so we denote the value of $\Pr(X_{i+k} = v | X_i = u)$ by $\Pr(u \to_k v)$, which is called the *k-step transition probability* from vertex u to vertex v. For all $k \geq 1$, $\Pr(u \to_k v)$ can be recursively formulated by

$$\Pr(u \to_k v) = \sum_{w \in I_G(v)} \Pr(u \to_{k-1} w) \Pr(w \to_1 v).$$

We can also formulate transition probabilities in the form of matrices. Suppose $V(G) = \{v_1, v_2, \ldots, v_n\}$. For $k \geq 1$, let $\mathbf{W}^{(k)}$ be the matrix of k-step transition probabilities, that is, $\mathbf{W}^{(k)}_{i,j} = \Pr(v_i \to_k v_j)$ for $1 \leq i, j \leq n$. We have $\mathbf{W}^{(1)} = \mathbf{A}$, where $\mathbf{A}$ is the *adjacency matrix* of G with rows normalized, that is, $\mathbf{A}_{i,j} = 1/|O_G(v_i)|$ if $(v_i, v_j) \in E(G)$, and $\mathbf{A}_{i,j} = 0$ otherwise. For $k > 1$, we have

$$\mathbf{W}^{(k)} = \mathbf{W}^{(k-1)} \mathbf{W}^{(1)} = \mathbf{A}^k.$$

SimRank. *SimRank* is a structural-context similarity measure for vertices in a directed graph [14]. It is designed based on the intuition that two vertices are similar if their in-neighbors are similar too. Let $s(u, v)$ be the *SimRank similarity* between vertices u and v in a directed graph G. $s(u, v)$ is defined by

$$s(u, v) = \begin{cases} 1 & \text{if } u = v,\\ \frac{c \sum_{u' \in I_G(u)} \sum_{v' \in I_G(v)} s(u', v')}{|I_G(u)||I_G(v)|} & \text{otherwise,}\end{cases} \tag{2}$$

where $0 < c < 1$ is called the *delay factor*. The system of Eq. (2) can be reformulated using a matrix equation. Let $V(G) = \{v_1, v_2, \ldots, v_n\}$ and $\mathbf{S}$ be a matrix with n rows and n columns, where $\mathbf{S}_{i,j} = s(v_i, v_j)$ for $1 \leq i, j \leq n$. Let $\mathbf{A}$ be the column-normalized adjacency matrix of graph G. We have

$$\mathbf{S} = c\mathbf{A}^T \mathbf{S} \mathbf{A} - \mathrm{diag}(c\mathbf{A}^T \mathbf{S} \mathbf{A}) + \mathbf{I},$$

where $\mathrm{diag}(\mathbf{X})$ is the diagonal matrix whose diagonal components are the diagonal components of $\mathbf{X}$. In many literatures [7], [19], [21], [39], $\mathbf{S}$ is often approximated as

$$\mathbf{S} = c\mathbf{A}^T \mathbf{S} \mathbf{A} + (1 - c)\mathbf{I}. \tag{3}$$

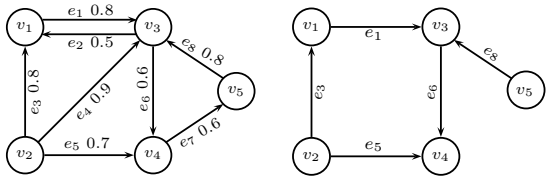

(a) Uncertain graph. (b) Possible world.

Fig. 1. An uncertain graph and one of its possible worlds.

Uncertain Graphs. An *uncertain graph* is a tuple (V, E, P), where V is a set of vertices, E is a set of arcs, and $P : E \to (0, 1]$ is a function assigning *existence probabilities* to the arcs. Particularly, $P(e)$ is the probability that arc e exists in practice. For clarity, we denote an uncertain graph by a written letter such as $\mathcal{G}$ and denote a deterministic graph by a printed letter such as G. Let $V(\mathcal{G})$, $E(\mathcal{G})$ and $P_{\mathcal{G}}$ be the vertex set, the edge set and the existence probability function of an uncertain graph $\mathcal{G}$, respectively. Let $I_{\mathcal{G}}(v)$ and $O_{\mathcal{G}}(v)$ be the sets of in-neighbors and out-neighbors of vertex v in an uncertain graph $\mathcal{G}$, respectively.

Under the *possible world semantics* [16], [18], [30], [44], [45], an uncertain graph $\mathcal{G}$ represents a probability distribution over all its *possible worlds*. More precisely, a *possible world* of $\mathcal{G}$ is a deterministic graph G such that $V(G) = V(\mathcal{G})$ and $E(G) \subseteq E(\mathcal{G})$. Let $\Omega(\mathcal{G})$ be the set of all possible worlds of $\mathcal{G}$ and $\mathcal{G} \Rightarrow G$ be the event that $\mathcal{G}$ exists in the form of its possible world G in practice. Following previous works [7], [16], [18], [30], [44]–[46], we reasonably assume that the existence probabilities of edges are mutually independent. Hence, the probability of event $\mathcal{G} \Rightarrow G$ is

$$\Pr(\mathcal{G} \Rightarrow G) = \prod_{e \in E(G)} P_{\mathcal{G}}(e) \cdot \prod_{e \in E(\mathcal{G}) \setminus E(G)} (1 - P_{\mathcal{G}}(e)). \tag{4}$$

It is easy to verify that $\sum_{G \in \Omega(\mathcal{G})} \Pr(\mathcal{G} \Rightarrow G) = 1$. Fig. 1 shows an uncertain graph $\mathcal{G}$ and one of its possible worlds G. For the possible world G in Fig. 1(b), we have $\Pr(\mathcal{G} \Rightarrow G) = P_{\mathcal{G}}(e_1) P_{\mathcal{G}}(e_3) P_{\mathcal{G}}(e_5) P_{\mathcal{G}}(e_6) P_{\mathcal{G}}(e_8)(1 - P_{\mathcal{G}}(e_2))(1 - P_{\mathcal{G}}(e_4))(1 - P_{\mathcal{G}}(e_7)) \approx 0.0043$.

III. Random Walks on Uncertain Graphs

In this section we give a formal definition of a random walk on an uncertain graph. Let $\mathcal{G}$ be an uncertain graph. Under the possible world model, $\mathcal{G}$ encodes a probability distribution over $\Omega(\mathcal{G})$, the set of all possible worlds of $\mathcal{G}$. Let $X_0, X_1, X_2, \ldots$ be a random walk on any possible

world of $\mathcal{G}$. For all $v_0, v_1, \ldots, v_n \in V(\mathcal{G})$, the probability $\Pr(X_n = v_n | X_0 = v_0, \ldots, X_{n-1} = v_{n-1})$ generally takes different values on different possible worlds of $\mathcal{G}$.

Let $\Pr_G(E)$ denote the probability of an event E on a possible world G, and let $\Pr_{\mathcal{G}}(E)$ denote the probability of an event E on a possible world of $\mathcal{G}$ selected at random according to the probability distribution given in Eq. (4). Then, we have

$$\begin{aligned}
&\Pr_{\mathcal{G}}(X_n = v_n | X_0 = v_0, \ldots, X_{n-1} = v_{n-1}) \\
=&\sum_{G \in \Omega(\mathcal{G})} \Pr_G(X_n = v_n | X_0 = v_0, \ldots, X_{n-1} = v_{n-1}) \Pr(\mathcal{G} \Rightarrow G) \\
=&\sum_{G \in \Omega(\mathcal{G})} \Pr_G(X_n = v_n | X_{n-1} = v_{n-1}) \Pr(\mathcal{G} \Rightarrow G) \\
=&\Pr_{\mathcal{G}}(X_n = v_n | X_{n-1} = v_{n-1}).
\end{aligned} \tag{5}$$

The second equality is due to Markov's property of a random walk on a deterministic graph. The above equation states that, on an uncertain graph, the probability that a random walk is on vertex v_n at time n is independent of all the previous vertices except the vertex it stays at time $n-1$.

We now define the *k-step transition probability* from a vertex u to a vertex v on uncertain graph $\mathcal{G}$ for $k \geq 1$. On a randomly chosen possible world of $\mathcal{G}$, the probability that a random walk is on vertex v at time n given that it is on vertex u at time $n-k$ is given by

$$\begin{aligned}
&\Pr_{\mathcal{G}}(X_n = v | X_{n-k} = u) \\
=&\sum_{G \in \Omega(\mathcal{G})} \Pr_G(X_n = v | X_{n-k} = u) \Pr(\mathcal{G} \Rightarrow G) \\
=&\sum_{G \in \Omega(\mathcal{G})} \Pr_G(u \to_k v) \Pr(\mathcal{G} \Rightarrow G).
\end{aligned} \tag{6}$$

Of course, $\Pr_{\mathcal{G}}(X_n = v | X_{n-k} = u)$ is fixed regardless of n, so we use $\Pr_{\mathcal{G}}(u \to_k v)$ to denote the value of $\Pr_{\mathcal{G}}(X_n = v | X_{n-k} = u)$ for any $n \geq k$.

For $k \geq 1$, let $\mathcal{W}^{(k)}$ be the matrix of k-step transition probabilities on uncertain graph $\mathcal{G}$, where $\mathcal{W}^{(k)}_{i,j} = \Pr_{\mathcal{G}}(v_i \to_k v_j)$. By Eq. (6), we have

$$\mathcal{W}^{(k)} = \sum_{G \in \Omega(\mathcal{G})} \Pr(\mathcal{G} \Rightarrow G) \mathbf{W}^{(k)}_G,$$

where $\mathbf{W}^{(k)}_G$ is the matrix of k-step transition probabilities on possible world G.

IV. Computing Transition Probabilities

In this section we propose an algorithm for computing the k-step transition probability, $\Pr_{\mathcal{G}}(u \to_k v)$, from a vertex u to a vertex v in an uncertain graph $\mathcal{G}$. By Eq. (6), we have

$$\begin{aligned}
&\Pr_{\mathcal{G}}(u \to_k v) = \Pr_{\mathcal{G}}(X_k = v | X_0 = u) \\
=&\sum_{v_1, v_2, \ldots, v_{k-1} \in V(\mathcal{G})} \Pr_{\mathcal{G}}(X_1 = v_1, X_2 = v_2, \ldots, X_k = v | X_0 = u).
\end{aligned} \tag{7}$$

In the above equation, if $u, v_1, v_2, \ldots, v_{k-1}, v$ is not a walk, then $\Pr_{\mathcal{G}}(X_1 = v_1, X_2 = v_2, \ldots, X_{k-1} = v_{k-1}, X_k = v | X_0 = u) = 0$. Hence, we only need to consider the vertices $v_1, v_2, \ldots, v_{k-1}$ such that $u, v_1, v_2, \ldots, v_{k-1}, v$ is a walk.

For convenience of presentation, let $v_0 = u$ and $v_k = v$. For any walk $W = v_0, v_1, \ldots, v_k$, we call $\Pr_{\mathcal{G}}(X_1 = v_1, X_2 = v_2, \ldots, X_k = v_k | X_0 = v_0)$ the *walk probability* of W given that W starts from v_0. According to the possible world model,

$$\begin{aligned}
&\Pr_{\mathcal{G}}(X_1 = v_1, X_2 = v_2, \ldots, X_k = v_k | X_0 = v_0) \\
=&\sum_{G \in \Omega(\mathcal{G})} \Pr_G(X_1 = v_1, X_2 = v_2, \ldots, X_k = v_k | X_0 = v_0) \Pr(\mathcal{G} \Rightarrow G) \\
=&\sum_{G \in \Omega(\mathcal{G})} \prod_{i=0}^{k-1} \Pr_G(X_{i+1} = v_{i+1} | X_i = v_i) \Pr(\mathcal{G} \Rightarrow G).
\end{aligned}$$

Hence, computing the k-step transition probability $\Pr_{\mathcal{G}}(u \to_k v)$ reduces to computing the walk probabilities of all walks starting from u and staying at v after k additional transitions.

Note that, on a deterministic graph G, the walk probability $\Pr_G(X_1 = v_1, X_2 = v_2, \ldots, X_k = v_k | X_0 = v_0)$ can be easily computed by Eq. (1), that is,

$$\Pr_G(X_1 = v_1, \ldots, X_k = v_k | X_0 = v_0) = \prod_{i=1}^{k} \Pr_G(v_{i-1} \to_1 v_i).$$

However, this simple method cannot be generalized to computing walk probabilities on an uncertain graph because

$$\begin{aligned}
&\Pr_{\mathcal{G}}(X_1 = v_1, X_2 = v_2, \ldots, X_k = v_k | X_0 = v_0) \\
=&\sum_{G \in \Omega(\mathcal{G})} \left(\Pr(\mathcal{G} \Rightarrow G) \prod_{i=0}^{k-1} \Pr_G(v_i \to_1 v_{i+1}) \right)
\end{aligned}$$

and

$$\prod_{i=0}^{k-1} \Pr_{\mathcal{G}}(v_i \to_1 v_{i+1}) = \prod_{i=0}^{k-1} \sum_{G \in \Omega(\mathcal{G})} \Pr(\mathcal{G} \Rightarrow G) \Pr_G(v_i \to_1 v_{i+1}).$$

The two equations above are generally *unequal* unless none of $v_0, v_1, \ldots, v_k$ are the same. Intuitively, if $v_i = v_j$ for some $1 \leq i, j \leq k$ and $i \neq j$, then, on any possible world G, the transition from v_i to v_{i+1} and the transition from v_j to v_{j+1} are not independent. This finding distinguishes our work from the work by Du et al. [7], in which they make an unreasonable assumption that $\Pr_{\mathcal{G}}(X_1 = v_1, \ldots, X_k = v_k | X_0 = v_0) = \prod_{i=1}^{k} \Pr_{\mathcal{G}}(v_{i-1} \to_1 v_i)$.

In the following we propose an algorithm for computing walk probabilities in Section IV-A and an algorithm for computing k-step transition probabilities in Section IV-B.

A. Computing Walk Probabilities

Let $W = v_0, v_1, \ldots, v_k$ be a walk on uncertain graph $\mathcal{G}$. Let $V(W)$ be the set of vertices in W. We have $|V(W)| \leq k+1$ because a vertex may appear multiple times in W. For every vertex $v \in V(W)$, let $O_W(v)$ be the set of out-neighbors of v in W, and let $c_W(v)$ be the number of occurrences of arcs from v to a vertex in W. We have $c_W(v) \geq |O_W(v)|$ because the walk may transit from v to a certain vertex in $O_W(v)$ multiple times.

For all possible worlds G of $\mathcal{G}$, if W is a walk in G, it follows from Eq. (1) that

$$\begin{aligned}
&\Pr_G(X_1 = v_1, X_2 = v_2, \ldots, X_k = v_k | X_0 = v_0) \\
=&\prod_{v \in V(W)} \mathrm{inv}(|O_G(v)|)^{c_W(v)},
\end{aligned}$$

where $\mathrm{inv}(x) = 1/x$ if $x \neq 0$, and $\mathrm{inv}(x) = 1$ otherwise. If W is not a walk in G, then $\Pr_G(X_1 = v_1, X_2 = v_2, \ldots, X_k = v_k | X_0 = v_0) = 0$. Therefore, the walk probability $\Pr_{\mathcal{G}}(X_1 = v_1, X_2 = v_2, \ldots, X_k = v_k | X_0 = v_0)$ can be computed by

$$\begin{aligned}
&\Pr_{\mathcal{G}}(X_1 = v_1, X_2 = v_2, \ldots, X_k = v_k | X_0 = v_0) \\
=&\sum_{G \in \Omega(\mathcal{G}), W \vdash G} \Pr(\mathcal{G} \Rightarrow G) \prod_{v \in V(W)} \mathrm{inv}(|O_G(v)|)^{c_W(v)},
\end{aligned} \tag{8}$$

where $W \vdash G$ represents that W is a walk in G.

Due to the independence assumption on the edges of $\mathcal{G}$, we have the following lemma, which gives an equivalent formulation of Eq. (8).

Lemma 1:

$$\Pr_{\mathcal{G}}(X_1 = v_1, \ldots, X_k = v_k | X_0 = v_0) = \prod_{v \in V(W)} \alpha_W(v), \quad (9)$$

where

$$\alpha_W(v) = \sum_G \Pr(\mathcal{G} \Rightarrow G) \text{inv}(|O_G(v)|)^{c_W(v)}.$$

The summation in the equation for $\alpha_W(v)$ is over all possible worlds G such that $(v, w) \in E(G)$ for all $w \in O_W(v)$.

Proof: First, we rewrite Eq. (8) as

$$\begin{aligned} &\Pr_{\mathcal{G}}(X_1 = v_1, X_2 = v_2, \ldots, X_k = v_k | X_0 = v_0) \\ =&\mathbf{E}[\Pr_G(X_1 = v_1, X_2 = v_2, \ldots, X_k = v_k | X_0 = v_0)] \\ =&\mathbf{E}[\prod_{v \in V(W)} \text{inv}(|O_G(v)|)^{c_W(v)}]. \end{aligned} \quad (10)$$

That is, the walk probability of W on $\mathcal{G}$ is the expectation of the walk probability over all the possible world graphs G of $\mathcal{G}$. Since the out edges of a vertex v is independent of the out edges of another vertex u. Thus the expectation and the product operations in Eq. (10) could be exchanged as

$$\begin{aligned} &\Pr_{\mathcal{G}}(X_1 = v_1, X_2 = v_2, \ldots, X_k = v_k | X_0 = v_0) \\ =&\mathbf{E}[\prod_{v \in V(W)} \text{inv}(|O_G(v)|)^{c_W(v)}] \\ =& \prod_{v \in V(W)} \mathbf{E}[\text{inv}(|O_G(v)|)^{c_W(v)}] \\ =& \prod_{v \in V(W)} (\sum_G \Pr(\mathcal{G} \Rightarrow G) \text{inv}(|O_G(v)|)^{c_W(v)}). \end{aligned}$$

Let $\alpha_W(v) = \sum_G \Pr(\mathcal{G} \Rightarrow G) \text{inv}(|O_G(v)|)^{c_W(v)}$. The lemma holds. ■

For all $v \in V(W)$, we can compute the term $\alpha_W(v)$ in Eq. (9) in polynomial time. The method is described as follows. Observe that, for each $v \in V(W)$, $c_W(v)$ is a constant; while $|O_G(v)|$ varies on different possible worlds G. Our method is based on evaluating the probability distribution of $|O_G(v)|$ across all possible worlds G of $\mathcal{G}$ such that $(v, w) \in E(G)$ for all $w \in O_W(v)$.

Let $O_{\mathcal{G}}(v) \setminus O_W(v) = \{w_1, w_2, \ldots, w_n\}$. For $0 \le j \le i \le n$, let $r(i, j)$ represent the probability that only j vertices in $w_1, w_2, \ldots w_i$ are connected to v in a randomly selected possible world of $\mathcal{G}$. Then, we have

$$\begin{aligned} &r(0,0) = 1, \\ &r(i,0) = r(i-1,0)(1 - P_{\mathcal{G}}((v,w_i))) && \text{for } 1 \le i \le n, \\ &r(i,i) = r(i-1,i-1)P_{\mathcal{G}}((v,w_i)) && \text{for } 1 \le i \le n, \\ &r(i,j) = r(i-1,j-1)P_{\mathcal{G}}((v,w_i)) \\ &\quad + r(i-1,j)(1 - P_{\mathcal{G}}((v,w_i))) && \text{for } 1 \le j < i \le n. \end{aligned}$$

Naturally, the probability that $|O_G(v)| = x$ in a randomly selected possible world G such that $(v, w) \in E(G)$ for all $w \in O_W(v)$ equals $r(n, x - |O_W(v)|)$. Thus, we have

$$\alpha_W(v) = \prod_{w \in O_W(v)} P_{\mathcal{G}}((v,w)) \sum_{x=0}^{n} r(n,x) \text{inv}(x + |O_W(v)|)^{c_W(v)}. \quad (11)$$

By Lemma 1 and Eq. (11), we immediately have the WalkPr algorithm as described in Figure 2 for computing the walk probability of a walk W in an uncertain graph $\mathcal{G}$. For every vertex $v \in V(W)$, lines 3–9 compute the values $r(i, j)$ for $0 \le j \le i \le n$ in a bottom-up manner, which totally run in $O((|O_{\mathcal{G}}(v)| - |O_W(v)|)^2)$ time. Using the values $r(n, x)$ for $0 \le x \le n$, line 10 computes $\alpha_W(v)$ in $O(|O_{\mathcal{G}}(v)|)$ time, and line 11 multiplies p by $\alpha_W(v)$. Thus, the running time of WalkPr is $\sum_{v \in V(W)} (|O_{\mathcal{G}}(v)| - |O_W(v)|)^2$. Let d be the average out-degree of the vertices in $\mathcal{G}$. The time complexity of WalkPr is therefore $O(|W|d^2)$.

Algorithm WalkPr($\mathcal{G}, W$)
1: $p \leftarrow 1$
2: **for all** $v \in V(W)$ **do**
3: //let $O_{\mathcal{G}}(v) \setminus O_W(v) = \{w_1, w_2, \ldots, w_n\}$
4: $r(0,0) \leftarrow 1$
5: **for** $i \leftarrow 1$ **to** n **do**
6: $r(i,0) \leftarrow r(i-1,0)(1 - P_{\mathcal{G}}((v,w_i)))$
7: $r(i,i) \leftarrow r(i-1,i-1)P_{\mathcal{G}}((v,w_i))$
8: **for** $j \leftarrow 1$ **to** $i-1$ **do**
9: $r(i,j) \leftarrow r(i-1,j-1)P_{\mathcal{G}}((v,w_i)) + r(i-1,j)(1 - P_{\mathcal{G}}(v,w_i)))$
10: $\alpha \leftarrow \prod_{w \in O_W(v)} P_{\mathcal{G}}((v,w)) \sum_{x=0}^{n} r(n,x) \text{inv}(x + |O_W(v)|)^{c_W(v)}$
11: $p \leftarrow p\alpha$
12: **return** p

Fig. 2. Algorithm WalkPr.

TABLE I. AN EXAMPLE OF COMPUTING WALK PROBABILITIES.

v	v_1	v_2	v_3	v_4
$O_W(v)$	$\{v_3\}$	$\{v_3\}$	$\{v_1, v_4\}$	$\{v_2\}$
$c_W(v)$	2	1	3	2
$O_{\mathcal{G}}(v) \setminus O_W(v)$	$\emptyset$	$\{v_1\}$	$\emptyset$	$\{v_5\}$
$r(n,x)$	$r(0,0) = 1$	$r(1,0) = 0.2$ $r(1,1) = 0.8$	$r(0,0) = 1$	$r(1,0) = 0.4$ $r(1,1) = 0.6$
$\alpha_W(v)$	0.64	0.54	0.0375	0.385

Example Consider the uncertain graph $\mathcal{G}$ illustrated in Fig. 1(a). We demonstrate how to compute the walk probability of a walk $W = v_1, v_3, v_1, v_3, v_4, v_2, v_3, v_4, v_2$ in $\mathcal{G}$ by Algorithm WalkPr. We have $V(W) = \{v_1, v_2, v_3, v_4\}$. For each $v \in V(W)$, we illustrate $O_W(v)$, $c_W(v)$, $O_{\mathcal{G}}(v) \setminus O_W(v)$, $r(n, x)$ and $\alpha_W(v)$ in Table I. Consequently, the output of WalkPr is $0.64 \times 0.54 \times 0.0375 \times 0.385 = 0.0049896$.

B. Computing k-step Transition Probabilities

We now propose the algorithm for computing the k-step transition probability, $\Pr_{\mathcal{G}}(u \to_k v)$, from a vertex u to a vertex v in an uncertain graph $\mathcal{G}$. By Eq. (7), $\Pr(u \to_k v)$ is the summation of walk probabilities of all walks starting from u and ending at v after k transitions. To further improve efficiency, instead of computing $\Pr(u \to_k v)$ from scratch, we compute $\Pr(u \to_k v)$ based on the $(k-1)$-step transition probabilities $\Pr(u \to_{k-1} w)$ for all vertices w such that (w, v) is an arc. Our incremental method is based on the following lemmas.

Lemma 2: Let $W = v_0, v_1, \ldots, v_k$ be a walk on $\mathcal{G}$ and $(v_k, v_{k+1}) \in E(\mathcal{G})$. Then, $W' = v_0, v_1, \ldots, v_{k+1}$ is also a walk on $\mathcal{G}$, and

$$\frac{\Pr_{\mathcal{G}}(X_1 = v_1, \ldots, X_{k+1} = v_{k+1} | X_0 = v_0)}{\Pr_{\mathcal{G}}(X_1 = v_1, \ldots, X_k = v_k | X_0 = v_0)} = \frac{\alpha_{W'}(v_k)}{\alpha_W(v_k)}.$$

Proof: By Lemma 1, we have

$$\Pr_{\mathcal{G}}(X_1 = v_1, \ldots, X_k = v_k | X_0 = v_0) = \prod_{v \in V(W)} \alpha_W(v)$$

and

$$\Pr_{\mathcal{G}}(X_1 = v_1, \ldots, X_k = v_k | X_0 = v_0) = \prod_{v \in V(W')} \alpha_{W'}(v).$$

Since W and W' only differs on the last vertex, only $\alpha_W(v)$ changes. Thus, we have $\alpha_W(v) = \alpha_{W'}(v)$ for all vertices $v \in V(W)$ and $v \neq v_k$. Hence, this lemma holds. ■

In some cases when W is not too long, the above lemma can be simplified as the following one.

Lemma 3: Let $W = v_0, v_1, \ldots, v_k$ and $W' = v_0, v_1, \ldots, v_{k+1}$ be two walks on $\mathcal{G}$, where (v_k, v_{k+1}) is an arc of $\mathcal{G}$. If the minimum length of the cycles in $\mathcal{G}$ is at least k, then we have

$$\frac{\Pr_{\mathcal{G}}(X_1 = v_1, \ldots, X_{k+1} = v_{k+1} | X_0 = v_0)}{\Pr_{\mathcal{G}}(X_1 = v_1, \ldots, X_k = v_k | X_0 = v_0)} = \Pr_{\mathcal{G}}(v_k \rightarrow_1 v_{k+1}).$$

Proof: Since the minimum length of the cycles in $\mathcal{G}$ is at least k, the out arcs of vertices $v_1, v_2, \ldots, v_k$ in W is at most 1. Thus, we have

$$\alpha_W(v_k) = \Pr_{\mathcal{G}}(v_k \rightarrow_1 v_{k+1}),$$

for $v_1, v_2, \ldots, v_k$ in W.

By Eq. (9), the lemma holds. ■

Based on Lemmas 2 and 3, we propose the TransPr algorithm to compute k-step transition probabilities as described in Figure 3. The input of TransPr is an uncertain graph $\mathcal{G}$ and an integer K. The output of TransPr is the k-step transition probability matrices $\mathcal{W}^{(1)}, \mathcal{W}^{(2)}, \ldots, \mathcal{W}^{(K)}$.

The algorithm first computes $\mathcal{W}^{(1)}$ and keeps it in main memory (line 1). The space used to store $\mathcal{W}^{(1)}$ is $O(|E(\mathcal{G})|)$. Then, we write $\mathcal{W}^{(1)}$ to the 1-step transition probability matrix file on disk. Particularly, for every 1-step walk $W = u, v$, we write to disk a tuple composed by the walk W, the walk probability of W, and the value $\alpha_W(v)$. Then, we compute the length ℓ of the shortest cycle in $\mathcal{G}$ using the algorithm proposed in [12] (line 2).

In the main loop (lines 3–18), we compute $\mathcal{W}^{(k+1)}$ based on $\mathcal{W}^{(k)}$ for $1 \leq k \leq K-1$. For each specific k, we scan the walk probability file of k-step walks on disk. For each tuple (W, p, α) in the file, W is a walk of length k, p is the walk probability of W, and α is the value $\alpha_W(v)$, where v is the last vertex in W. For every vertex $w \in O_{\mathcal{G}}(v)$, we append w to the end of W and thus obtain a new walk W' of length $k+1$. We compute the walk probability of W' based on the walk probability of W either by Lemma 3 (line 9) or by Lemma 2 (line 12). Moreover, we compute the value $\alpha_{W'}(w)$. After this, we write W', the walk probability of W' and the value $\alpha_{W'}(w)$ to the walk probability file of $(k+1)$-step walks on disk (line 14). When the scanning over the walk probability file of k-step walks is completed, we sort the tuples (W', p', α') in the walk probability file of $(k+1)$-step walks according to the start and the end vertices of W'. For all walks with the same start vertex u and the same end vertex v, we compute the $(k+1)$-step transition probability $\Pr_{\mathcal{G}}(u \rightarrow_{k+1} v)$ by summing up the walk probabilities of all these walks (line 17). Finally, we write the tuple $(u, v, \Pr_{\mathcal{G}}(u \rightarrow_{k+1} v))$ to the file of the $(k+1)$-step transition probability matrix $\mathcal{W}^{(k+1)}$ (line 18).

Algorithm TransPr
Input: an uncertain graph $\mathcal{G}$ and an integer K
Output: $\mathcal{W}^{(1)}, \mathcal{W}^{(2)}, \ldots, \mathcal{W}^{(K)}$
1: compute $\mathcal{W}^{(1)}$ and write it to disk
2: $\ell \leftarrow$ the length of the shortest cycle in $\mathcal{G}$
3: **for** $k \leftarrow 1$ **to** $K-1$ **do**
4: **for** every tuple (W, p, α) in the walk probability file of k-step walks **do**
5: $v \leftarrow$ the last vertex in W
6: **for all** $w \in O_{\mathcal{G}}(v)$ **do**
7: $W' \leftarrow$ the walk obtained by appending v to W
8: **if** $k \leq \ell$ **then**
9: $p' \leftarrow p\mathcal{W}^{(1)}_{vw}$
10: $\alpha' \leftarrow 1$
11: **else**
12: $\alpha' \leftarrow \alpha_{W'}(w)$
13: $p' \leftarrow p\alpha'/\alpha$
14: write (W', p', α') to the walk probability file of $(k+1)$-step walks
15: sort all tuples (W, p, α) according to the start and the end vertices of W
16: **for** each pair of vertices u and v **do**
17: $\Pr_{\mathcal{G}}(u \rightarrow_{k+1} v) \leftarrow$ summation of walk probabilities of all walks starting at u and ending at v
18: write $(u, v, \Pr_{\mathcal{G}}(u \rightarrow_{k+1} v))$ to the file of $\mathcal{W}^{(k+1)}$

Fig. 3. Algorithm TransPr.

V. SimRank Similarities on Uncertain Graphs

In this section we give a formal definition of SimRank similarity on an uncertain graph. First, let us review a random-walk-based definition of SimRank similarity on a deterministic graph. Let G be a deterministic graph and $\mathbf{A}$ the adjacency matrix of G with columns normalized. We have the following definition of the *kth SimRank similarity matrix* $\mathbf{S}^{(k)}$.

$$\mathbf{S}^{(0)} = \mathbf{I},$$
$$\mathbf{S}^{(n)} = c\mathbf{A}^T\mathbf{S}^{(n-1)}\mathbf{A} + (1-c)\mathbf{I} \qquad \text{for } n > 0,$$

where $\mathbf{I}$ is the identity matrix, and $0 < c < 1$ is the delay factor. Since G is a deterministic graph, we have $\mathbf{A}^k = \mathbf{W}^{(k)}$ for $k > 0$. Moreover, $(\mathbf{A}^T)^k = (\mathbf{A}^k)^T$. Let $\mathbf{W}^{(0)} = \mathbf{I}$. Thus,

$$\mathbf{S}^{(n)} = c^n(\mathbf{W}^{(n)})^T\mathbf{W}^{(n)} + (1-c)\sum_{k=0}^{n-1} c^k(\mathbf{W}^{(k)})^T\mathbf{W}^{(k)}.$$

Note that the element at the ith row and the jth column of $(\mathbf{W}^{(k)})^T\mathbf{W}^{(k)}$ is the probability of two random walks starting from vertices i and j, respectively, meeting at the same vertex after exactly k transitions. The theorem below states that $\mathbf{S}^{(n)}$ converges to the *SimRank similarity matrix* $\mathbf{S}$ as $n \rightarrow +\infty$.

Theorem 1: $\lim_{n\rightarrow+\infty} \mathbf{S}^{(n)} = \mathbf{S}$.

Based on the possible worlds model of uncertain graphs, we define SimRank similarity on an uncertain graph as follows.

Definition 1: For a uncertain graph $\mathcal{G}$, a delay factor $0 < c < 1$, and $n \geq 1$, the *nth SimRank similarity* between two vertices u and v in $\mathcal{G}$, denoted by $s^{(n)}_{\mathcal{G}}(u, v)$, is defined by

$$\begin{aligned} s^{(n)}_{\mathcal{G}}(u, v) = c^n \sum_{w \in V(\mathcal{G})} \Pr_{\mathcal{G}}(u \rightarrow_n w) \Pr_{\mathcal{G}}(v \rightarrow_n w) \\ + (1-c)\sum_{k=0}^{n-1} c^k \Pr_{\mathcal{G}}(u \rightarrow_k w) \Pr_{\mathcal{G}}(v \rightarrow_k w). \end{aligned} \tag{12}$$

The *SimRank similarity* between u and v, denoted by $s_{\mathcal{G}}(u, v)$, is defined by $s_{\mathcal{G}}(u, v) = \lim_{n\rightarrow+\infty} s^{(n)}_{\mathcal{G}}(u, v)$.

The following theorem gives an upper bound on the error between $s^{(n)}_{\mathcal{G}}(u, v)$ and $s_{\mathcal{G}}(u, v)$. The proof is omitted.

Theorem 2: For $n \geq 1$, $|s^{(n)}_{\mathcal{G}}(u, v) - s_{\mathcal{G}}(u, v)| \leq c^{n+1}$.

As n increases, the error between $s^{(n)}_{\mathcal{G}}(u, v)$ and $s_{\mathcal{G}}(u, v)$ decreases exponentially. Therefore, similar to SimRank computation on a deterministic graph, we can use $s^{(n)}_{\mathcal{G}}(u, v)$ as a good approximation of $s_{\mathcal{G}}(u, v)$ when n is sufficiently large.

The following theorem shows that the SimRank similarity defined on uncertain graphs is a *generalization* of the SimRank similarity defined on deterministic graphs.

Theorem 3: Let G be a deterministic graph and $\mathcal{G}$ be an uncertain graph with $V(\mathcal{G}) = V(G)$, $E(\mathcal{G}) = E(G)$, and $P_{\mathcal{G}}(e) = 1$ for all $e \in E(\mathcal{G})$. For all $u, v \in V(\mathcal{G})$, the SimRank similarity between u and v on $\mathcal{G}$ equals the SimRank similarity between u and v on G.

VI. Computing SimRank Similarities

In this section we propose several algorithms for computing the SimRank similarity between two vertices in an uncertain graph, namely the baseline algorithm in Section VI-A, the sampling algorithm in Section VI-B, the two-phase algorithm in Section VI-C and the speeding-up algorithm in Section VI-D.

A. Baseline Algorithm

We first describe the baseline algorithm. Given as input an uncertain graph $\mathcal{G}$, two vertices u, v, a real number $c \in (0,1)$ and an integer $n > 0$, we first compute the transition probability matrices $\mathcal{W}^{(1)}, \mathcal{W}^{(2)}, \ldots, \mathcal{W}^{(n)}$ by the TransPr algorithm. Then, we compute $s^{(n)}_{\mathcal{G}}(u,v)$ by Eq. (12) and return $s^{(n)}_{\mathcal{G}}(u,v)$ as an approximation of $s_{\mathcal{G}}(u,v)$.

To compute the term $\sum_{w \in V(\mathcal{G})} \Pr_{\mathcal{G}}(u \rightarrow_k w) \Pr_{\mathcal{G}}(v \rightarrow_k w)$ in Eq. (12), we need to read the two columns of $\mathcal{W}^{(k)}$ corresponding to u and v, respectively. Since $\mathcal{W}^{(k)}$ is not sparse, we store $\mathcal{W}^{(k)}$ in external memory. To facilitate data access, we store the elements of $\mathcal{W}^{(k)}$ column-by-column in consecutive blocks on disk. Let B be the size of a disk block. Reading a column requires $O(|V(\mathcal{G})|/B)$ I/O's. Hence, the total number of I/O's of the baseline algorithm is $O(n|V(\mathcal{G})|/B)$.

B. Sampling Algorithm

The second algorithm for computing $s^{(n)}_{\mathcal{G}}(u,v)$ is based on random sampling. In this algorithm, we estimate each term $\sum_{w \in V(\mathcal{G})} \Pr_{\mathcal{G}}(u \rightarrow_k w) \Pr_{\mathcal{G}}(v \rightarrow_k w)$ in Eq. (12) via sampling. For $k > 0$, the sampling procedure is as follows.

Let N be a sufficiently large integer. We randomly sample N walks $W^u_1, W^u_2, \ldots, W^u_N$ starting from u and N walks $W^v_1, W^v_2, \ldots, W^v_N$ starting from v, all of which are of length n. Each walk W should be sampled with the walk probability of W. A simple method to do this is to first sample a possible world G of $\mathcal{G}$ with probability $\Pr(\mathcal{G} \Rightarrow G)$ and then randomly select a walk W on G. Indeed, this method is inefficient. In our algorithm, we adopt a more efficient method. Without loss of generality, let us take u as the starting vertex. For every vertex of $\mathcal{G}$, we record the status if the vertex has been visited by W. Initially, W only contains u. Then, we extend W by iteratively appending vertices at the end of W until the length of W reaches k. In each iteration, we check whether the last vertex w in W has been visited. If w has not been visited, we first instantiate every arc e leaving w with probability $P_{\mathcal{G}}(e)$ and designate w as visited. If w has been visited, we select uniformly at random an out-neighbor z of w via an instantiated arc leaving w, and append z at the end of W.

After sampling N walks $W^u_1, W^u_2, \ldots, W^u_N$ starting from u and N walks $W^v_1, W^v_2, \ldots, W^v_N$ starting from v, we estimate $\sum_{w \in V(\mathcal{G})} \Pr_{\mathcal{G}}(u \rightarrow_k w) \Pr_{\mathcal{G}}(v \rightarrow_k w)$ for all $k > 0$. For ease of presentation, let us denote $\sum_{w \in V(\mathcal{G})} \Pr_{\mathcal{G}}(u \rightarrow_k w) \Pr_{\mathcal{G}}(v \rightarrow_k w)$ by $m^{(k)}(u,v)$. For $1 \le i \le N$ and $1 \le j \le n$, let $I(i,j) = 1$ if the jth vertex in W^u_i and the jth vertex in W^v_i are the same; otherwise, $I(i,j) = 0$. Hence, $m^{(k)}(u,v)$ can be estimated by

$$\widehat{m}^{(k)}(u,v) = \frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{n} I(i,j). \tag{13}$$

By Chernoff's bound, we have the following lemma on the error of $\widehat{m}^{(k)}(u,v)$.

Algorithm Sampling($\mathcal{G}$, u, v, n, N)
1: **for** $i \leftarrow 1$ **to** N **do**
2: $W^u_i \leftarrow u$
3: **for** $j \leftarrow 1$ **to** n **do**
4: // let w be the last vertex of W^u_i
5: **if** w is not visited by W^u_i **then**
6: sample all arcs leaving w according to its existing probability
7: mark w visited by W^u_i
8: choose an instantiated neighbor x of w at random
9: append x onto W^u_i
10: **for** $i \leftarrow 1$ **to** N **do**
11: $W^v_i \leftarrow v$
12: **for** $j \leftarrow 1$ **to** n **do**
13: // let w be the last vertex of W^v_i
14: **if** w is not visited by W^v_i **then**
15: sample all arcs leaving w according to its existing probability
16: mark w visited by W^v_i
17: choose an instantiated out-neighbor z of w at random
18: append z at the end of W^v_i
19: **for** $k \leftarrow 1$ **to** n **do**
20: compute $\widehat{m}^{(k)}(u,v)$ according to Eq. (13)
21: compute $\widehat{s}^{(n)}_{\mathcal{G}}(u,v)$ Eq. (14)
22: **return** $\widehat{s}^{(n)}_{\mathcal{G}}(u,v)$

Fig. 4. Algorithm Sampling.

Lemma 4: For $\epsilon > 0$ and $0 < \delta < 1$, if $N \ge \frac{3}{\epsilon^2} \ln \frac{2}{\delta}$, then

$$\Pr(|m^{(k)}(u,v) - \widehat{m}^{(k)}(u,v)| \le \epsilon) \ge 1 - \delta.$$

Consequently, by Eq. (12), we can estimate $s^{(n)}_{\mathcal{G}}(u,v)$ by

$$\widehat{s}^{(n)}_{\mathcal{G}}(u,v) = c^n \widehat{m}^{(n)}(u,v) + (1-c) \sum_{k=0}^{n-1} c^k \widehat{m}^{(k)}(u,v). \tag{14}$$

The Sampling algorithm in Fig. 4 illustrates the details of our sampling algorithm. We have the following result on the approximation error of Sampling.

Theorem 4: For $\epsilon > 0$ and $0 < \delta < 1$, if $N \ge \frac{3}{\epsilon^2} \ln \frac{2}{\delta}$,

$$\Pr\left(|s^{(n)}_{\mathcal{G}}(u,v) - \widehat{s}^{(n)}_{\mathcal{G}}(u,v)| \le (c - c^n)\epsilon\right) \ge 1 - \delta.$$

The time complexity of Sampling is $O(Nnd)$, where d is the average degree of the vertices in $\mathcal{G}$. The Sampling algorithm is more efficient than the baseline algorithm because it performs less I/O's and uses less memory.

C. The Two-stage Algorithm

To take the advantages of both the baseline algorithm and the sampling algorithm, we propose the two-stage algorithm. The algorithm is based on the following observations.

(1) When k is small, the number of nonzero elements in the transition probability matrix $\mathcal{W}^{(k)}$ is far less than $|V(\mathcal{G})|^2$. Especially, there are only $|E(\mathcal{G})|$ nonzero elements in $\mathcal{W}^{(1)}$. It decreases the number of I/O's to read the columns of $\mathcal{W}^{(k)}$. Thus, the method used in the baseline algorithm is efficient in computing the exact value of $m^{(k)}(u,v)$.

(2) When k is large, the error of the estimated value $\widehat{m}^{(k)}(u,v)$ computed by the sampling method is less than $c^k\epsilon$ with probability at least $1 - \delta$.

Inspired by the two observations, we compute $s^{(n)}_{\mathcal{G}}(u,v)$ in two stages. Let $1 < l < n$.

Stage 1. For $1 \le k \le l$, compute $m^{(k)}(u,v)$ using the exact method given in the baseline algorithm.

Stage 2. For $l < k \le n$, estimate $m^{(k)}(u,v)$ using the method given in the Sampling algorithm. Let $\widehat{m}^{(k)}(u,v)$ be the estimated value of $m^{(k)}(u,v)$.

Algorithm Speedup($\mathcal{G}$, u, v, n, N)
1: $U^{(0)} \leftarrow \{u\}$, $V^{(0)} \leftarrow \{v\}$
2: insert $(0, \mathbf{1})$ into $\mathbf{M}_u$ and $\mathbf{M}'_v$
3: **for** $k \leftarrow 0$ **to** $n-1$ **do**
4: $\quad U^{(k+1)} \leftarrow \emptyset$
5: $\quad$ **for** each vertex w in $U^{(k)}$ **do**
6: $\quad\quad$ **for** each vertex x in $O_{\mathcal{G}}(w)$ **do**
7: $\quad\quad\quad \mathbf{M}_x[k+1] = \mathbf{M}_x[k+1] \vee (\mathbf{M}_w[k] \wedge \mathbf{F}_{(}w, x))$
8: $\quad\quad\quad$ **if** $\mathbf{M}_x[k+1] \neq \mathbf{0}$ **then**
9: $\quad\quad\quad\quad$ insert x into $U^{(k+1)}$
10: **for** $k \leftarrow 0$ **to** $n-1$ **do**
11: $\quad V^{(k+1)} \leftarrow \emptyset$
12: $\quad$ **for** each vertex w in $V^{(k)}$ **do**
13: $\quad\quad$ **for** each vertex x in $O_{\mathcal{G}}(w)$ **do**
14: $\quad\quad\quad \mathbf{M}'_x[k+1] = \mathbf{M}'_x[k+1] \vee (\mathbf{M}'_w[k] \wedge \mathbf{F}_{(}w, x))$
15: $\quad\quad\quad$ **if** $\mathbf{M}'_x[k+1] \neq \mathbf{0}$ **then**
16: $\quad\quad\quad\quad$ insert x into $V^{(k+1)}$
17: **for** $k \leftarrow 1$ **to** n **do**
18: $\quad$ compute $\widehat{m}^{(k)}(u, v)$ according to Eq. (16)
19: compute $\widehat{s}^{(n)}_{\mathcal{G}}(u, v)$ according to Eq. (15)
20: **return** $\widehat{s}^{(n)}_{\mathcal{G}}(u, v)$

Fig. 5. Algorithm Speedup.

After the above two stages, we estimate $s^{(n)}_{\mathcal{G}}(u, v)$ by

$$\widehat{s}^{(n)}_{\mathcal{G}}(u, v) = c^n \widehat{m}^{(n)}(u, v) + (1-c) \sum_{k=l+1}^{n-1} c^k \widehat{m}^{(k)}(u, v) + (1-c) \sum_{k=0}^{l} c^k m^{(k)}(u, v). \tag{15}$$

By Theorem 4, we immediately have the corollary below.

Corollary 1: For $\epsilon > 0$ and $0 < \delta < 1$, if $N \geq \frac{3}{\epsilon^2} \ln \frac{2}{\delta}$, $\Pr\left(|s^{(n)}_{\mathcal{G}}(u, v) - \widehat{s}^{(n)}_{\mathcal{G}}(u, v)| \leq (c^{l+1} - c^n)\epsilon\right) \geq 1 - \delta$.

By tuning l, we can make a tradeoff between the time efficiency and the approximation error. By Corollary 1, the relative error of the output of the two-stage algorithm is bounded by $(c^{l+1} - c^n)\epsilon / s^{(n)}_{\mathcal{G}}(u, v)$. When l is larger, the relative error of $s^{(n)}_{\mathcal{G}}(u, v)$ decreases exponentially. Meanwhile, the execution time of the two-stage algorithm increases because we need to read more components of the matrices on disk. If l is carefully chosen such that $\mathcal{W}^{(1)}, \mathcal{W}^{(2)}, \ldots, \mathcal{W}^{(l)}$ fit into main memory, the two-stage algorithm is as efficient as the Sampling algorithm. For example, let $l = 1$, $\mathcal{W}^{(1)}$ only consumes $O(|E|)$ space. In this setting, for an SimRank similarity value which is about $c/10$, the relative approximation error is near $c\epsilon$, which is an order of magnitude better than the Sampling algorithm.

D. Speeding-up Technique

We now propose the technique for speeding up the sampling process in the sampling algorithm and the two-phase algorithm. Remember that, given two vertices u and v, we sample independently at random N walks $W^u_1, W^u_2, \ldots, W^u_N$ starting from u and N walks $W^v_1, W^v_2, \ldots, W^v_N$ starting from v, all of which are of length n. As analyzed in Section VI-B, the expected time to sample each of these walks is $O(nd)$, where d is the average degree of the vertices of $\mathcal{G}$. To reduce the total time of sampling all these walks, we utilize an observation that $W^u_1, W^u_2, \ldots, W^u_N$ (or $W^v_1, W^v_2, \ldots, W^v_N$) usually have significant overlaps, which can be used to reduce redundant extensions of walks. For example, Fig. 6 shows ten walks starting from v_1 and v_2, respectively, sampled from the uncertain graph shown in Fig. 1(a). Notice that v_3 occurs five times in the first step of the five walks starting from v_1, and v_4 occurs twice in the first step of the fives walks starting from v_2. If we could compress the traverses of these extensions of walks, a lot of redundant operations could be reduced, which will further speed up the algorithm.

	Walks starting from v_1		Walks starting from v_2
W_1	$v_1, v_3, v_1, v_3, v_1, v_3$	W_6	$v_2, v_1, v_3, v_4, v_5, v_3$
W_2	$v_1, v_3, v_4, v_5, v_3, v_4$	W_7	$v_2, v_3, v_1, v_3, v_4, v_5$
W_3	$v_1, v_3, v_4, v_5, v_3, v_1$	W_8	$v_2, v_4, v_5, v_3, v_1, v_3$
W_4	$v_1, v_3, v_1, v_3, v_4, v_5$	W_9	$v_2, v_3, v_4, v_5, v_3, v_1$
W_5	$v_1, v_3, v_4, v_5, v_3, v_4$	W_{10}	$v_2, v_4, v_5, v_3, v_1, v_3$

Fig. 6. An example of ten walks of length 5 sampled at random.

In our speeding-up method, we associate every vertex w of $\mathcal{G}$ with a hash table, denoted by $\mathbf{M}_w$. We call $\mathbf{M}_w$ the *counting table* of w. The key of each entry of the hash table is an integer. For each key k, the value corresponding to k is a N-dimensional bit vector, denoted by $\mathbf{M}_w[k]$. For $1 \leq i \leq N$, the ith bit of $\mathbf{M}_w[k]$ is 1 if and only if w is the kth vertex in the walk W^u_i. If the hash table $\mathbf{M}_w$ does not contain key k, then we conceptually have $\mathbf{M}_w[k] = \mathbf{0}$. Similarly, we associate every vertex w of $\mathcal{G}$ with a hash table $\mathbf{M}'_w$. For each key k, the ith bit of $\mathbf{M}'_w[k]$ is 1 if and only if w is the kth vertex in the walk W^v_i, where $1 \leq i \leq N$. Thus, the information of the walks $W^u_1, W^u_2, \ldots, W^u_N$ (or $W^v_1, W^v_2, \ldots, W^v_N$) is *losslessly* encoded in all nonempty hash tables $\mathbf{M}_w$ (or $\mathbf{M}'_w$).

Moreover, we associate every arc $e = (w, x)$ of $\mathcal{G}$ with a N-dimensional bit vector $\mathbf{F}_e$, called the *filter vector* of e. We construct the filter vectors of all the arcs of $\mathcal{G}$ offline. Initially, we set the bit vectors of all arcs to $\mathbf{0}$, where $\mathbf{0}$ is a vector with all bits set to 0. For all vertices w of $\mathcal{G}$ and all $1 \leq i \leq N$, we first instantiate every arc $e = (w, x)$ with probability $P_{\mathcal{G}}(e)$, where $x \in O_{\mathcal{G}}(w)$. Then, we select one instantiated arc e leaving w uniformly at random and set the ith bit of the filter vector $\mathbf{F}_e$ to 1.

We now describe our method for speeding up the sampling process. Suppose that the filter vectors of all edges have been constructed offline. In our new method, to obtain the sampled walks $W^u_1, W^u_2, \ldots, W^u_N$, we need not to perform the sampling process N times. Instead, we start from vertex u and perform N sampling processes simultaneously by leveraging the common substructures among samples. The Speedup algorithm in Figure 5 describes the process of the speedup algorithm. The details of the method are given as follows.

Step 1. Let $U^{(k)}$ be the set of vertices that are probable to be visited at the kth step of a walk starting from u. Initially, we set $U^{(0)} = \{u\}$ and $U^{(k)} = \emptyset$ for all $1 \leq k \leq n$. Since u is the 0th vertex in all sampled walks, we insert an entry $(0, \mathbf{1})$ to the counting table $\mathbf{M}_u$, where $\mathbf{1}$ is the bit vector with all bits set to 1. For k from 0 to $n-1$, we perform the kth iteration as follows. For each vertex $w \in U^{(k)}$, we retrieve the set, $O_{\mathcal{G}}(w)$, of all out-neighbors of w. Each out-neighbor x of w is probable to be visited at the $(k+1)$-th step of a walk starting from u. Recall that the bit vector $\mathbf{M}_w[k]$ records that in which sampled walks, w is the kth vertex. Suppose the ith bit of $\mathbf{M}_w[k]$ is 1, that is, w is the kth vertex in the ith sampled walk W^u_i. For every out-neighbor x of w, if the ith bit of the filter vector $\mathbf{F}_{(w,x)}$ is 1, that is, the walk goes from w to x in W^u_i, then x is certainly the $(k+1)$-th vertex in W^u_i, that is, the ith bit of $\mathbf{M}_x[k+1]$ should be set to 1. Thus, we update $\mathbf{M}_x[k+1]$ to $\mathbf{M}_x[k+1] \vee (\mathbf{M}_w[k] \wedge \mathbf{F}_{(w,x)})$, where $\vee$ and $\wedge$ denote the bit-wise OR and the bit-wise AND operations, respectively. If $\mathbf{M}_x[k+1] \neq \mathbf{0}$, we insert x to $U^{(k+1)}$.

TABLE II. SUMMARY OF DATASETS USED IN EXPERIMENTS.

Dataset	Number of vertices	Number of edges
PPI1	2708	7123
PPI2	2369	249080
PPI3	19247	17096006
Condmat	31163	240058
Net	1588	5484
DBLP	1560640	8517894

Step 2. Let $V^{(k)}$ be the set of vertices that are probable to be visited at the kth step of a walk starting from u. Similar to Step 1, we start from vertex v and perform another n iterations to compute $V^{(k)}$ for $1 \leq k \leq n$ and update the counting tables $\mathbf{M}'_w$ associated with the vertices.

Step 3. We compute $m^{(k)}(u, v)$ based on the counting tables $\mathbf{M}_w$ and $\mathbf{M}'_w$. In particular, for $0 \leq k \leq n$, the value of $m^{(k)}(u, v)$ can be estimated by

$$\widehat{m}^{(k)}(u, v) = \frac{1}{N} \sum_{w \in U^{(k)} \cap V^{(k)}} \|\mathbf{M}_w[k] \wedge \mathbf{M}'_w[k]\|_1, \quad (16)$$

where $\|\mathbf{x}\|$ denotes the 1-norm of a vector $\mathbf{x}$, that is, the number of 1's in $\mathbf{x}$. It is easy to verify that the value $\widehat{m}^{(k)}(u, v)$ computed by the above equation is the same as the one computed by Eq. (13). Consequently, we can estimate $s^{(n)}_{\mathcal{G}}(u, v)$ either by Eq. (14) or by Eq. (15).

VII. EXPERIMENTS

We conducted extensive experiments to evaluate the effectiveness of our SimRank similarity measure and the performance of the proposed algorithms. We present the experimental results in this section.

A. Experimental Setting

We implemented the proposed algorithms in C++, including the baseline algorithm (Baseline), the sampling algorithm (Sampling), the two-stage algorithm (SR-TS) and the two-stage algorithm with the speed-up technique (SR-SP). All experiments were run on a Linux machine with 2.5GHz Intel Core i5 CPU and 16GB of RAM.

Table II summarizes the datasets used in our experiments. *PPI1*, *PPI2* and *PPI3* are three protein-protein interaction networks extracted from [18] and the STRING[1] database. *Condmat* and *Net* are two widely used co-authorship networks provided by [29]. *DBLP* is a co-authorship obtained from AMiner[2]. For *Condmat*, *Net* and *DBLP*, we set the uncertainty of each edge using the method in [44].

If not otherwise stated, on each uncertain graph, we ran the algorithms 1000 times. For each time, we selected a pair of vertices uniformly at random from the input uncertain graph and computed the SimRank similarity between the vertices. We evaluated the average execution time and relative error of the 1000 runs. Therefore, the execution time and the relative error reported in this section are actually the average execution time and the average relative error, respectively. By default, we set $n = 5$, $c = 0.6$ and $N = 1000$. As reported in the following, the SimRank similarity generally converges within 5 iterations.

[1] http://string-db.org

[2] http://aminer.org/citation

TABLE III. DIFFERENCES BETWEEN SimRank-I AND OTHER MEASURES.

Dataset	Similarity	Avg. Bias	Max. Bias	Min. Bias
Net	SimRank-II	0.048	0.219	0
	SimRank-III	0.039	0.323	0
	Jaccard-I	0.072	1	2.16×10^{-7}
	Jaccard-II	0.160	0.770	0
PPI1	SimRank-II	0.075	0.609	0
	SimRank-III	0.031	0.214	0
	Jaccard-I	0.130	1	0
	Jaccard-II	0.143	0.913	0

B. Experimental Results

Differences between Similarity Measures. First, we test the effectiveness of our SimRank similarity measure by comparing the similarities computed using our SimRank similarity measure and those computed using other similarity measures. We choose 1000 pairs of vertices on *Net* and *PPI1* uniformly at random. For each pair of vertices, we compute their SimRank similarity by the Baseline algorithm (SimRank-I) and compare it with the similarities computed by other methods, including the SimRank similarity computed on the deterministic graph obtained by removing uncertainty from the uncertain graph (SimRank-II), the SimRank similarity computed by Du et al.'s algorithm [7] (SimRank-III), the Jaccard similarity computed by the algorithm in [44] (Jaccard-I) and the Jaccard similarity computed on the deterministic graph obtained by removing uncertainty from the uncertain graph (Jaccard-II).

Fig. 7 reports the differences between SimRank-I and the other similarities. Fig. 7(a) and Fig. 7(f) show the SimRank-I similarities of 1000 randomly selected vertex pairs on *Net* and *PPI1*, respectively. The vertex pairs are sorted in decreasing order of their SimRank-I similarities. Fig. 7(b)–7(e) compare SimRank-I with SimRank-II, SimRank-III, Jaccard-I and Jaccard-II computed on *Net*, respectively. Fig. 7(g)–7(j) compare SimRank-I with SimRank-II, SimRank-III, Jaccard-I and Jaccard-II computed on *PPI1*, respectively. Note that all similarities are normalized to within $[0, 1]$. The differences are summarized in Table III. We observe that (1) when SimRank-I varies slightly, the other similarities may vary significantly; (2) when SimRank-I decreases, the other similarities may increase. The differences between SimRank-I and Jaccard-I and Jaccard-II are most significant because the Jaccard similarity cannot measure the similarity between vertices without common neighbors. SimRank-I is different from SimRank-II since SimRank-II does not consider uncertainty in graphs. SimRank-I also differs from SimRank-III since SimRank-III is based on an unreasonable assumption as we mentioned in Section IV.

Convergence. In this experiment, we examined the convergence of our SimRank computation algorithms. We varied the number n of iterations from 1 to 10 and computed the SimRank similarities of 1000 randomly selected vertex pairs by the Baseline algorithm. Fig. 8 shows the effects of n on the average and the maximum SimRank similarities between these 1000 vertex pairs on *PPI1*, *PPI2*, *Net* and *Condmat*. Obviously, the SimRank similarities remain stable after 5 iterations. This experiment verifies that the approximated SimRank similarity converges to the exact similarity as n becomes larger.

Efficiency. In this experiment, we compared the execution time of the algorithms Baseline, Sampling, SR-TS and SR-SP. Fig. 9 illustrates the average execution time of the algorithms.

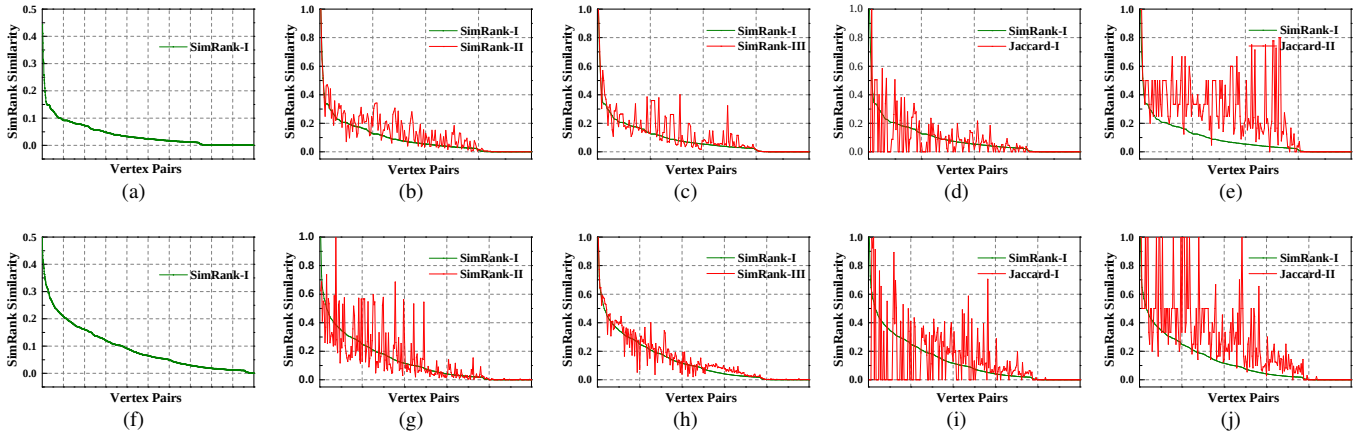

Fig. 7. Differences between similarity measures. (a) SimRank-I on *Net*. (b) SimRank-I vs. SimRank-II on *Net*. (c) SimRank-I vs. SimRank-III on *Net*. (d) SimRank-I vs. Jaccard-I on *Net*. (e) SimRank-I vs. Jaccard-II on *Net*. (f) SimRank-I on *PPI1*. (g) SimRank-I vs. SimRank-II on *PPI1*. (h) SimRank-I vs. SimRank-III on *PPI1*. (i) SimRank-I vs. Jaccard-I on *PPI1*. (j) SimRank-I vs. Jaccard-II on *PPI1*.

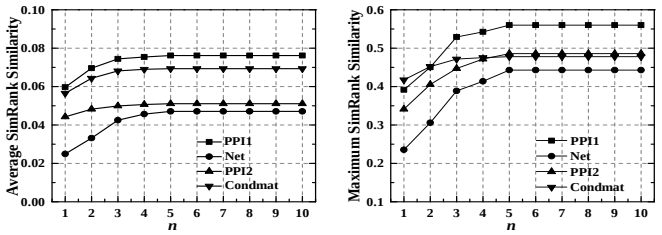

(a) Avg. SimRank similarity vs. n. (b) Max. SimRank similarity vs. n.

Fig. 8. Effects of the number of iterations on SimRank similarity.

For SR-TS and SR-SP, we set $l = 1, 2, 3$, respectively. From Fig. 9, we have the following observations.

1) Baseline is faster than Sampling and SR-TS on *PPI2*, *PPI3* and *Condmat* but is much slower on *DBLP*. This is because *PPI2*, *PPI3* and *Condmat* are small graphs, so their transition probability matrices can fit into main memory. However, each of the transition probability matrices $\mathcal{W}_2, \mathcal{W}_3, \ldots, \mathcal{W}_5$ of *DBLP* cannot fit into main memory. Thus, Baseline incurs high I/O cost on *DBLP*.

2) SR-SP is much faster than SR-TS on all the datasets. The speedup on *PPI2* and *PPI3* are more than 30 and 15, respectively. This is because SR-SP uses the speed-up technique, and the bit-wise operations to reduce sampling time.

3) The execution time of Sampling is independent of the input graph size. Sampling is faster on *DBLP* than on *PPI2* because the time complexity of Sampling is only related to the number of sampled walks and the density of the input graph. The execution time of Sampling on *PPI3* is very high because *PPI3* is very dense.

4) The execution time of SR-TS is comparable to Sampling. The parameter l of the two-phase algorithm has little effect on the execution time of SR-TS because the execution time of SR-TS is dominated by the sampling process.

Accuracy. We examined the accuracy of the algorithms by running them on all datasets 1000 times and computing the average relative error of the 1000 runs. Since it is hard to compute the exact SimRank similarity between two vertices, we take the SimRank similarity s^* computed by Baseline as the baseline and compute the relative error by $|s - s^*|/s^*$, where s is the similarity computed by a tested algorithm.

Fig. 10 shows the relative errors of the algorithms. We have two observations. 1) The relative errors of the algorithms are very small. In particular, the relative error of Sampling is about 10%, and the relative errors of SR-TS and SR-SP are nearly 1%. 2) The relative errors of SR-TS and SR-SP decrease with the growth of parameter l. This is consistent with Lemma 4. For the same l, the relative error of SR-TS and SR-SP are comparable. Since SR-TS is much more efficient than SR-SP, SR-SP is superior to SR-TS in practice.

Effects of Parameter N. In this experiment, we tested the effects of the number N of sampled walks on the efficiency and the accuracy of the algorithms. Since the execution time of Sampling is comparable to SR-TS, we only tested SR-TS and SR-SP in our experiment. We set $l = 1$.

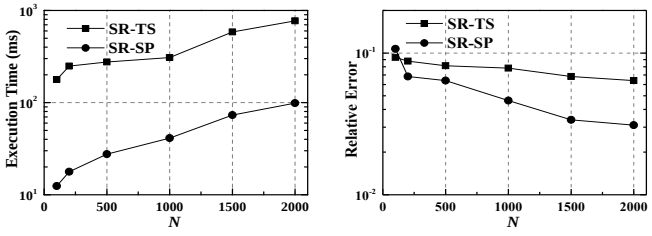

(a) Execution time vs. N. (b) Relative error vs. N.

Fig. 11. Effects of parameter N on execution time and relative error.

Fig. 11 shows the execution time and the relative error of SR-TS and SR-SP with respect to N on *Condmat*. We observe that the execution time of SR-TS and SR-SP grow sub-linearly with respect to N, and the relative errors decrease with the growth of N. When N is sufficient large, the relative error varies slightly. We can observe that the relative error is less than 5% for both algorithms when $N \geq 1000$. The experimental results on the other datasets are similar.

Scalability. In the last experiment, we tested the scalability of the algorithms on synthetic uncertain graphs. We generated a series of uncertain graphs with 2M vertices and 2M–10M

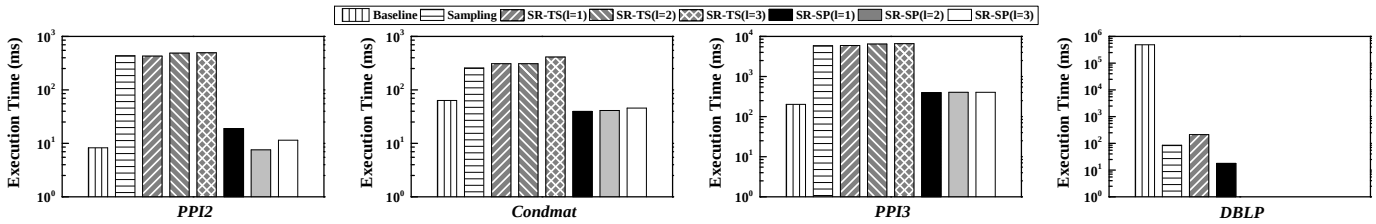

Fig. 9. Execution time of algorithms Baseline, Sampling, SR-TS and SR-SP. For SR-TS and SR-SP, we set $l = 1, 2, 3$, respectively.

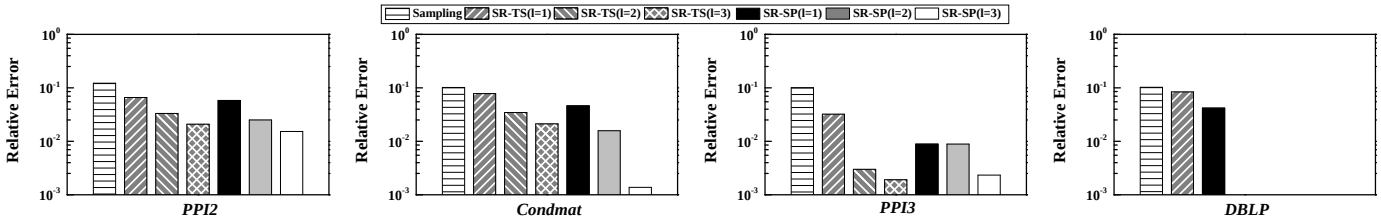

Fig. 10. Relative error of algorithms Sampling, SR-TS and SR-SP. For SR-TS and SR-SP, we set $l = 1, 2, 3$, respectively.

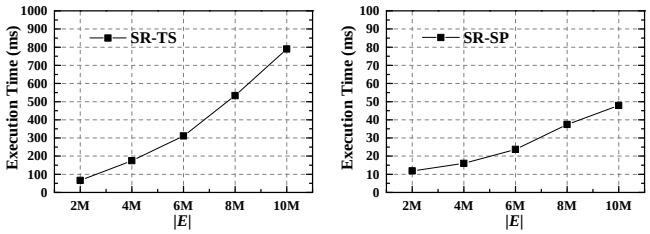

Fig. 12. Scalability of SR-TS and SR-SP with respect to graph size.

edges. The structures of the uncertain graphs were generated using the R-MAT model [5], and the probabilities of the edges were generated uniformly at random within $[0, 1]$. Let $N = 1000$, $n = 5$ and $l = 1$. We ran SR-TS and SR-SP on each uncertain graph 1000 times. Fig. 12 illustrates the average execution time of SR-TS and SR-SP. We observe that the execution time of both SR-TS and SR-SP grow almost linear with the number of edges. This is because the density of the graph is proportional to the number of edges, and the time complexity of SR-TS and SR-SP are highly dependent on the density of the graph. This experimental results show that our proposed algorithms attain high scalability.

C. Case Study

We demonstrate two case studies to show the effectiveness of our SimRank similarity measure. One is the detection of proteins with similar biological functions in a protein-protein interaction (PPI) network. The other is the aggregation of similar objects in graph-based entity resolution algorithms.

Detecting Similar Proteins. In this case study, we use our SimRank similarity measure on uncertain graphs to find similar proteins in a PPI network. Here we find the top-20 similar protein pairs by two methods. The first one is the SimRank measure proposed in this paper (USIM), and the second one is the SimRank measure without considering the uncertainties (DSIM) in the PPI network. Fig. 13 report the top-20 similar

(CTK2,CTK1) (MRPL19, MRPL11) (YHR197W, IPI3)
(YPL166W,CIS1) (YDR279W,YLR154C) (NUP85,SEH1)
(NUP85,NUP120) (SEH1,NUP120) (BZZ1,VRP1) (PSF1,CTF4)
(CTF8,CTF18) (RRN11,RRN7) (BUL1,YHR131C)
(NUP85,NUP84) (SEH1,NUP84) (NUP120,NUP84)
(PHO85,PHO80) **(UBP13,UBP9)** (ESC8 IOC3) (CDC53 GRR1)

(a) Top-20 similar protein pairs by USIM

(YPR143W YKL014C) (YMR067C DOA1) **(YPL166W,CIS1)**
(VIP1 YBR267W) **(YHR197W, IPI3) (NUP85,SEH1)**
(IPL1 YDR415C) (GRH1 SFB2) (LTV1 LYS9) (TUB4 SPC72)
(SAR1 MSK1) **(YDR489W CTF4)** (MBP1 YKR077W)
(ESC8 IOC3) (NUF2 SPC25) **(UBP13,UBP9)** (PHO85,PHO80)
(MCK1 RIC1) (VPS28 YGR206W) (YKL014C SSF2)

(b) Top-20 similar protein pairs by DSIM

Fig. 13. Top-20 similar Proteins detected by USIM and DSIM.

protein pairs on the *PPI1* dataset. Here we use the MIPS[3] database as the ground truth, which provides many known protein complexes. Protein pairs within a common protein complex are thought to coordinate with each other in biological functions [41], [46]. The protein pairs in Fig. 13 that are contained in the same protein complex are marked in boldface.

Notice that 16 pairs of proteins in the top-20 results by USIM are contained in the same protein complex and the top-9 pairs are all in the same protein complex. Whereas, only 6 pairs of proteins in the top-20 results by DSIM are verified to be in the same protein complex and only 3 pairs in the top-10 results are in the same protein complex. This comparison results show that our SimRank similarity measure is capable of capturing the structural-context similarity between objects with inherent uncertainties such as the PPI network, which verify the effectiveness of our SimRank measure on uncertain graphs once again.

Also, in Fig. 14 we report the top-5 similar proteins with respect to a specific protein BUB1, which plays an important role in the mitotic process of a cell. The relationship of BUB1 and RGA1 has been examined in [27], which claims to find a novel pathway in the mitotic exit where BUB1 coordinates

[3] http://mips.helmholtz-muenchen.de

RGA1, SRC1, ISM1, YOL098C, SHE4

Fig. 14. Top-5 Similar protein with respect to protein BUB1.

with RGA1 on the spindle.

Entity Resolution. In this case study, we apply the SimRank similarity measure proposed in this paper to graph-based entity resolution (ER). Following the framework of the EIF algorithm [22], we develop two new algorithms. The first one is SimER that regards the entity graph as an uncertain graph and adopts the SimRank similarity measure proposed in this paper. The second one is SimDER that regards the entity graph as an deterministic graph and adopts the SimRank similarity measure on deterministic graphs. We compare the SimER algorithm and the SimDER algorithm with the EIF algorithm [22] and the DISTINCT algorithm [35] on the *DBLP* dataset. The experimental results are as follows.

First we compare the efficiency of the SimER algorithm, the SimDER algorithm and the EIF algorithm by varying the records size from 2000 to 5000. In the SimER algorithm and the SimDER algorithm, we set the similarity threshold for aggregating data records to be 0.1. We set the sampling size N to 1000 and use the speed up techniques in the implementation of SimER and SimDER. Fig. 15 reports the execution time of the three algorithms. Here the execution time of both the three algorithms all increases approximately linear to the record size because they follow the same framework but only be different on the similarity measures, which are also verified in [22]. The execution time of the EIF algorithm is a bit faster than the SimER algorithm and the SimDER algorithm. However, the variance is not significant. On average, the EIF algorithm is about 20% faster than the SimDER algorithm and 25% faster than the SimER algorithm. And the DISTINCT algorithm is about 25% faster than the SimDER algorithm and 30% faster than the SimER algorithm.

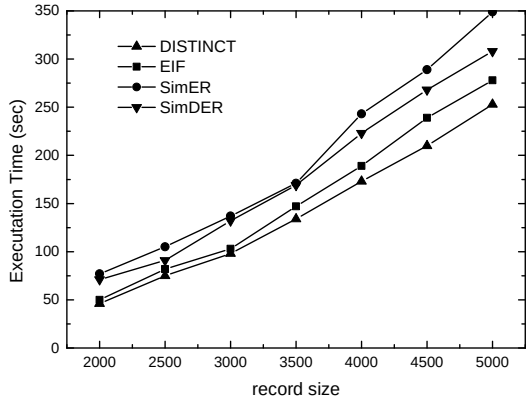

Fig. 15. Execution time of EIF, SimDER and SimER vs. record size.

Next we report the effectiveness of the proposed algorithms. We choose the same 8 representative author names in Table IV in the *DBLP* dataset, where each name corresponds to multiple authors (#authors) and multiple records (#records). We compare the *precision*, *recall* and F_1*-measure* of SimER, SimDER, EIF and DISTINCT. The experimental results are reported in Table V.

TABLE IV. 8 NAMES CORRESPONDING TO MULTIPLE RECORDS

Name	#authors	# ref	Name	#authors	# ref
Hui Fang	3	9	Ajay Gupta	4	16
Rakesh Kumar	2	38	Micheal Wagner	5	24
Bing Liu	6	11	Jim Smith	3	19
Wei Wang	14	177	Wei Wang	14	177

The *precision* of SimER and SimDER are comparable to that of EIF and is slightly better than that of DISTINCT. The recall of SimER and SimDER are much higher than that of EIF and DISTINCT. Also, SimER outperforms SimDER, EIF and DISTINCT in terms of *F-measure*. It verifies that the SimRank similarity is an effective method to measure the object similarities in graph data. More over, our SimRank similarity measure which take uncertainties into consideration are more effective for graphs with inherent uncertainties such as the entity resolution application.

VIII. Related Work

A. SimRank on Deterministic Graphs

SimRank [14] is a measure of similarities between vertices. A large number of algorithms have been proposed to compute SimRank similarities [8], [10], [19]–[21], [24], [31], [33], [36], [37], [39], [40]. These algorithms can be classified into three categories.

Iterative Algorithms. Jeh and Widom [14] proposed the first iterative algorithm for computing SimRank similarities by matrix computations. It computes the SimRank similarities between all pairs of vertices in $O(nd^2|V|^2)$ time, where n is the number of iterations, d is the average degree of vertices, and $|V|$ is the number of vertices. Later, [24] improved the running time to $O(nd|V|^2)$ by sum memorization, and [40] speeded up the algorithm by fast matrix multiplication. Further, [36] decreased the time complexity to $O(nd'|V|^2)$ by fine-grained memorization, where $d' < d$. For the single-pair SimRank computation problem, [23] computes the SimRank similarity between a pair of vertices in $O(nd^2 \min\{|V|^2, d^n\})$ time, and [11] further improved the running time to $O(n|E|^2 - |E|)$ by employing position probabilities, where $|E|$ is the number of edges. All these iterative algorithms require $O(|V|^2)$ space and thus can not scale to large graphs. Recently, [39] developed a space-efficient algorithm to compute the SimRank similarities between a subset of vertices in $O(|E| + n|V|)$ space.

Non-iterative Algorithms. In [21], a non-iterative SimRank matrix formula has been established based on the Kronecker product and singular vector decomposition (SVD). It precomputes several auxiliary matrices offline in $O(r^4|V|^2)$ time and then retrieves the SimRank similarities between a given vertex and all other vertices online in $O(r^4|V|)$ time, where r is the rank of the adjacency matrix of the graph. Moreover, this algorithm is capable of processing dynamic graphs. Later, [10] and [37] improved the running time to $O(r|V|)$ by Sylvester's equation. Recently, [25] proposed a novel approach to converting the SimRank formula into a linear equation. The algorithm requires only $O(|E|)$ space and computes all-pairs SimRank similarities and single-pair SimRank similarities in $O(n|V||E|)$ time and $O(n|E|)$ time, respectively.

Random-walk-based Algorithms. The SimRank similarity between two vertices u and v can be represented in form of the probabilities that two random walks starting from u and v,

TABLE V. COMPARISONS OF THE EXPERIMENTAL RESULTS OF SimER, SimDER, EIF AND DISTINCT ON MULTIPLE AUTHOR NAMES.

Name	SimER			SimDER			EIF			DISTINCT		
	Precision	*Recall*	*F-Measure*	*Precision*	*Recall*	*F-Measure*	*Precision*	*Recall*	*F-Measure*	*Precision*	*Recall*	*F-Measure*
Hui Fang	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Ajay Gupta	1.0	1.0	1.0	1.0	0.927	0.962	1.0	0.882	0.937	1.0	1.0	1.0
Rakesh Kumar	1.0	1.0	1.0	0.958	0.973	0.965	1.0	1.0	1.0	1.0	1.0	1.0
Micheal Wagner	0.993	0.880	0.933	1.0	0.859	0.924	1.0	0.620	0.765	1.0	0.395	0.566
Bing Liu	1.0	1.0	1.0	0.982	1.0	0.991	1.0	1.0	1.0	1.0	0.825	0.904
Jim Smith	0.988	0.963	0.975	0.973	0.978	0.975	1.0	0.810	0.895	0.888	0.926	0.907
Wei Wang	1.0	1.0	1.0	0.993	1.0	0.996	1.0	0.933	0.965	0.855	0.814	0.834
Bin Yu	0.980	0.975	0.989	0.989	0.752	0.854	0.977	0.595	0.746	1.0	0.658	0.794
Average	**0.995**	**0.977**	**0.987**	**0.987**	**0.936**	**0.958**	**1.0**	**0.855**	**0.914**	**0.968**	**0.827**	**0.876**

respectively, meet at the same vertex at the same time. The first random-walk-based algorithm [8] stores the fingerprints of N random walks as an index in $O(N|V|)$ space and computes single-pair SimRank similarities based on this index. In [19], a sampling algorithm was developed to compute the SimRank similarity between a single pair of vertices based on the random-walk interpretation of the linear formula given in [25]. For top-k SimRank computations, the algorithm [20] finds the k vertices with the highest SimRank similarities with respect to a given vertex. The studies in [31] improve the performance by compactly storing some samples of the origin graph. Different from this, [33] studied the problem of finding the top-k pairs of vertices with the highest SimRank similarities.

B. Random Walks on Uncertain Graphs

Although the concept of random walks on uncertain graphs has ever been used earlier in the literature [7], [17], [30], it is totally different from our definition in this paper. In [7], [17], [30], for a random walk that is on vertex u at time t, they sample the neighbors of u, randomly select a sampled neighbor v, and transit the state from u to v at time $t+1$. Therefore, for a vertex u and two different time t and t', the possible vertices that the walk can transit to from vertex u at time t are different from those that the walk can transit to from vertex u at time t'. However, on each possible world of the uncertain graph, the set of possible vertices that the walk can transit to from vertex u are the same all the time. Thus, the random walk [7], [17], [30] does not satisfy Markov's Property. In fact, our definition of random walks on uncertain graphs is the first one that satisfies Markov's Property (see Section III).

C. Similarities for Uncertain Graphs

Besides the algorithms introduced above, a number of variants of SimRank have also been investigated. [1], [15], [32], [38], [42]. However, almost all the studies on SimRank overlook uncertainty in graphs. In recent years, many studies on managing and mining uncertain graphs [16], [18], [45], [46] suggested that the effects of uncertainty on the quality of query or mining results have been undervalued. To the best of our knowledge, the only work on the SimRank on uncertain graphs has only been carried out in [7]. However, this work made an unreasonable assumption that the k-step transition probability matrix $\mathcal{W}^{(k)}$ equals to the kth power of the one-step transition probability matrix $\mathcal{W}^{(1)}$, i.e., $\mathcal{W}^{(k)} = (\mathcal{W}^{(1)})^k$ for all $k \geq 1$. This assumption is inconsistent with the possible world model of uncertain graphs that has been widely adopted in the literature.

Besides the SimRank similarity measure, three structural-context similarity measures between vertices of an uncertain graph have been proposed [44], namely the expected Jaccard similarity, the expected Dice similarity and the expected cosine similarity. Unlike the SimRank similarity, these three measures are only applicable to the vertices with common neighbors.

IX. CONCLUSIONS

This paper proposes the concepts and the algorithms related to the SimRank similarity on uncertain graphs. Unlike the most related work [7], our concepts are completely based on the possible world model of uncertain graphs. We propose three algorithms and a speeding-up technique for SimRank computation on an uncertain graph. The experimental results show that the algorithms are effective and efficient. To lay foundations for SimRank, we also study random walks on uncertain graphs for the first time. We reveal the critical differences between random walks on uncertain graphs and the counterparts on deterministic graphs.

ACKNOWLEDGEMENTS

We thank anonymous reviewers for their very useful comments and suggestions.

This work was partially supported by the 973 Program of China under grant No. 2011CB036202 and the National Natural Science Foundation of China under grant No. 61173023 and No.61532015.

REFERENCES

[1] I. Antonellis, H. G. Molina, and C. C. Chang. Simrank++: query rewriting through link analysis of the click graph. *PVLDB*, pages 408–421, 2008.

[2] R. Baeza-Yates, B. Ribeiro-Neto, et al. *Modern information retrieval*, volume 463. ACM press New York, 1999.

[3] A. A. Benczúr, K. Csalogány, and T. Sarlós. Link-based similarity search to fight web spam. In *AIRWEB*, 2006.

[4] I. Bhattacharya and L. Getoor. Entity resolution in graphs. *Mining graph data*, 2006.

[5] D. Chakrabarti, Y. Zhan, and C. Faloutsos. R-mat: A recursive model for graph mining. In *SDM*, pages 442–446, 2004.

[6] L. R. Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, 1945.

[7] L. Du, C. Li, H. Chen, L. Tan, and Y. Zhang. Probabilistic simrank computation over uncertain graphs. *Information Sciences*, 295:521–535, 2015.

[8] D. Fogaras and B. Rácz. Scaling link-based similarity search. In *WWW*, pages 641–650, 2005.

[9] F. Fouss, A. Pirotte, J.-M. Renders, and M. Saerens. Random-walk computation of similarities between nodes of a graph with application to collaborative recommendation. *IEEE Trans. on Knowledge and Data Engineering*, 19(3):355–369, 2007.

[10] Y. Fujiwara, M. Nakatsuji, H. Shiokawa, and M. Onizuka. Efficient search algorithm for simrank. In *ICDE*, pages 589–600, 2013.

[11] J. He, H. Liu, J. X. Yu, P. Li, W. He, and X. Du. Assessing single-pair similarity over graphs by aggregating first-meeting probabilities. *Information System*, 42:107–122, 2014.

[12] J. D. Horton. A polynomial-time algorithm to find the shortest cycle basis of a graph. *SIAM Journal on Computing*, 16(2):358–366, 1987.

[13] P. Jaccard. *Etude comparative de la distribution florale dans une portion des Alpes et du Jura*. Impr. Corbaz, 1901.

[14] G. Jeh and J. Widom. Simrank: a measure of structural-context similarity. In *KDD*, pages 538–543, 2002.

[15] R. Jin, V. E. Lee, and H. Hong. Axiomatic ranking of network role similarity. In *KDD*, pages 922–930, 2011.

[16] R. Jin, L. Liu, and C. C. Aggarwal. Discovering highly reliable subgraphs in uncertain graphs. In *KDD*, pages 992–1000, 2011.

[17] R. Jin, L. Liu, B. Ding, and H. Wang. Distance-constraint reachability computation in uncertain graphs. *PVLDB*, 4(9):551–562, 2011.

[18] G. Kollios, M. Potamias, and E. Terzi. Clustering large probabilistic graphs. *IEEE Trans. on Knowledge and Data Engineering*, 25(2):325–336, 2013.

[19] M. Kusumoto, T. Maehara, and K.-i. Kawarabayashi. Scalable similarity search for simrank. In *SIGMOD*, pages 325–336, 2014.

[20] P. Lee, L. V. Lakshmanan, and J. X. Yu. On top-k structural similarity search. In *ICDE*, pages 774–785, 2012.

[21] C. Li, J. Han, G. He, X. Jin, Y. Sun, Y. Yu, and T. Wu. Fast computation of simrank for static and dynamic information networks. In *EDBT*, pages 465–476, 2010.

[22] L. Li, H. Wang, H. Gao, and J. Li. Eif: A framework of effective entity identification. In *WAIM*, pages 717–728, 2010.

[23] P. Li, H. Liu, J. X. Yu, J. He, and X. Du. Fast single-pair simrank computation. In *SDM*, pages 571–582, 2010.

[24] D. Lizorkin, P. Velikhov, M. Grinev, and D. Turdakov. Accuracy estimate and optimization techniques for simrank computation. *PVLDB*, pages 422–433, 2008.

[25] T. Maehara, M. Kusumoto, and K.-i. Kawarabayashi. Scalable simrank join algorithm. In *ICDE*, pages 603–614, 2015.

[26] S. Melnik, H. Garcia-Molina, and E. Rahm. Similarity flooding: A versatile graph matching algorithm and its application to schema matching. In *ICDE*, pages 117–128, 2002.

[27] S. A. Nelson and J. A. Cooper. A novel pathway that coordinates mitotic exit with spindle position. *Molecular biology of the cell*, 18(9):3440–3450, 2007.

[28] M. E. Newman. Detecting community structure in networks. *European Physical Journal B-Condensed Matter and Complex System*, 38(2):321–330, 2004.

[29] M. E. Newman. Finding community structure in networks using the eigenvectors of matrices. *Physical review E*, 74(3):036104, 2006.

[30] M. Potamias, F. Bonchi, A. Gionis, and G. Kollios. K-nearest neighbors in uncertain graphs. *PVLDB*, pages 997–1008, 2010.

[31] Y. Shao, B. Cui, L. Chen, M. Liu, and X. Xie. An efficient similarity search framework for simrank over large dynamic graphs. *PVLDB*, pages 838–849, 2015.

[32] L. Sun, C. Cheng, X. Li, D. Cheung, and J. Han. On link-based similarity join. *PVLDB*, 2011.

[33] W. Tao, M. Yu, and G. Li. Efficient top-k simrank-based similarity join. *PVLDB*, pages 317–328, 2014.

[34] K. Whalen, B. Sadkhin, D. Davidson, and J. Gerlt. Sequence similarity networks for the protein universe. *The FASEB Journal*, 29(1 Supplement):573–17, 2015.

[35] X. Yin, J. Han, and P. S. Yu. Object distinction: Distinguishing objects with identical names. In *ICDE*, pages 1242–1246, 2007.

[36] W. Yu, X. Lin, and W. Zhang. Towards efficient simrank computation on large networks. In *ICDE*, pages 601–612, 2013.

[37] W. Yu, X. Lin, and W. Zhang. Fast incremental simrank on link-evolving graphs. In *ICDE*, pages 304–315, 2014.

[38] W. Yu, X. Lin, W. Zhang, L. Chang, and J. Pei. More is simpler: Effectively and efficiently assessing node-pair similarities based on hyperlinks. *PVLDB*, pages 13–24, 2013.

[39] W. Yu and J. A. McCann. Efficient partial-pairs simrank search on large networks. *PVLDB*, pages 569–580, 2015.

[40] W. Yu, W. Zhang, X. Lin, Q. Zhang, and J. Le. A space and time efficient algorithm for simrank computation. *WWW*, pages 327–353, 2012.

[41] B. Zhao, J. Wang, M. Li, F.-X. Wu, and Y. Pan. Detecting protein complexes based on uncertain graph model. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 11(3):486–497, 2014.

[42] P. Zhao, J. Han, and Y. Sun. P-rank: a comprehensive structural similarity measure over information networks. In *CIKM*, pages 553–562, 2009.

[43] W. Zheng, L. Zou, Y. Feng, L. Chen, and D. Zhao. Efficient simrank-based similarity join over large graphs. *PVLDB*, pages 493–504, 2013.

[44] Z. Zou and J. Li. Structural-context similarities for uncertain graphs. In *ICDM*, pages 1325–1330, 2013.

[45] Z. Zou, J. Li, H. Gao, and S. Zhang. Frequent subgraph pattern mining on uncertain graph data. In *CIKM*, pages 583–592, 2009.

[46] Z. Zou, J. Li, H. Gao, and S. Zhang. Finding top-k maximal cliques in an uncertain graph. In *ICDE*, pages 649–652, 2010.